SCIENCE FOR EVERYONE

V. I. VARSHAVSKY
D. A. POSPELOV

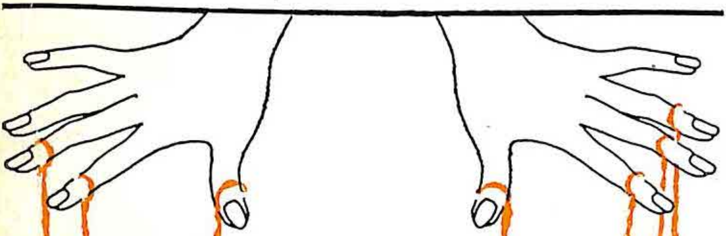

PUPPETS WITHOUT STRINGS

MIR

Science
for Everyone

В. И. Варшавский
Д. А. Поспелов

Оркестр играет без дирижера

Размышления об эволюции некоторых технических систем и управлении ими

«Наука» Москва

V. I. Varshavsky
D. A. Pospelov

Puppets Without Strings

Reflections on the Evolution
and Control of Some Man-Made Systems

Mir
Publishers
Moscow

Translated from Russian
by A. Kandaurov

First published 1988
Revised from the 1984 Russian edition

На английском языке
Printed in the Union of Soviet Socialist Republics

ISBN 5-03-000272-3

Contents

Instead of a Preface

The first public performance of the Persymphense (the name is an acronym stemming from what may be translated from Russian as the First Symphony Ensemble of the Moscow Soviet) in Moscow on 13 February 1922 was a sensation, for it amazed and puzzled both professional musicians and music lovers.

The reason was that the orchestra performed without a conductor. It was not as if the compositions they played were simple orchestral pieces. Instead the programme contained The Third (Heroical) Symphony and a violin concerto by Beethoven. The flow of exquisite sounds was running with so much harmony that the professional musicians in the audience left the concert at a complete loss. It seemed to them that there must have been some trick which they failed to notice, that there was an invisible somebody secretly conducting the orchestra and producing that unique impression which, of all things, only a conductor's will can make. For nobody but a conductor can offer a profoundly individual interpretation of a musical piece, impose a desired tempo, synchronize the variety of parts played by the different musical instruments and make a large orchestra sound in harmony. It is for this reason that all the performers on the stage are arranged so that they all

can see the conductor and follow his directions.

The arrangement of the Persymphense performers, however, was unusual. The strings were seated in a closed circle, some of them with their backs to the audience, while the wind-instruments were in the centre of the circle. Each musician could see the others, for each of the orchestra's performers had to listen to and to be listened to by all the rest. There was no trick. By interacting directly with each other, the brilliant performers of the Persymphense orchestra could dispense with a conductor.

For a whole decade the concerts of the orchestra were vastly popular both with the professional and general public. Everyone sought to unravel the mystery of this unusual group. Given a common purpose, that is their first-rate artistic interpretation of the music, each player did his best to achieve his own local aim by relying on his professional skill. Another illustration may be a traditional jazz band. To sum up, the success of the Persymphense orchestra was a triumph of decentralization over the centralized control exercised by a conductor. The decentralized control naturally evolved due to the joint effort of cooperating musicians. The situation could not be properly explained at the time because it was an enigma to the axioms of conventional logic.

It should be pointed out that complex processes due to local interactions between

their components instead of centralized control occur in both nature and human society much more frequently than it may seem at first sight. Consequently, to reveal how the decentralized control arises in such situations due to the collective interaction of the system's components is a far harder job than unravelling the mystery of the Persymphense orchestra. Apart from other things, this book is an attempt to cope with this subject.

We seek to present a popular account of the control problems that arise in complex systems which are more generally called large-scale systems in control theory. In systems of this kind, centralized control often gives way to decentralized control, the transition being a penalty for the system's complexity. This is because the system's complexity makes centralized control either inefficient or impracticable. How do large-scale systems arise, and is it possible that the category of large-scale system is merely a far-fetched nothing? We have tried to show in this book that large-scale man-made systems which surround us are steadily becoming more numerous and still more complex. The evolution of man-made systems out of the already existing ones goes on in much the same way as living organisms evolve. Decentralized control is but a natural product of this evolution. We hope to convince our readers that it is just so.

Chapter 1

Decentralized Control: the Whys and Wherefores

> A voice: "Before something appeared, nothing had existed."
>
> *Isidor Shtok*

1.1. The Man-Made World

When our primeval ancestor took a stone and tried for the first time to use it to give shape to a shapeless piece of rock, there appeared a realm of things that had previously not existed in nature. The things produced by human hands and brains make up what may now be called a *technocenosis*. Like a biological community which is called a biocenosis, a technocenosis is a set of different machines, instruments, systems and devices which can be grouped together by close and superficially unusual relationships. As inhabitants of a man-made world, we seldom pay attention to these relationships. Few of us notice that a door-handle is fixed a certain distance from the floor or that we have to apply a certain effort to open or close the door itself. None of us are amazed that a wood screw may be screwed in and out with what is known as screw-driver. You would not surprise anyone by telling him that a freight container placed on a train in Moscow may cross Europe over rivers by bridges and through

mountains by tunnels, then cross the Channel by ferry to find itself in London. Like a biocenosis, a technocenosis makes all the "organisms" it is composed of live according to the rules imposed by the whole community. If an aircraft-maker seeks to enrich a technocenosis with a new plane whose takeoff and landing run demands a landing strip longer than the one available in today's airfields, this plane can exist only in his imagination or the manufacturer's shop. The reason is simple: the plane does not have an "ecological niche" in the technocenosis related to flying.

In this book, we will often resort to examples from biology and use the terms applied to the structure and function of biological communities. The drawing of such parallels is not a whim on our part. We are convinced that the organic world created by nature and the technical world, which has been and still is being created by man, are very much alike. The similarity is far more deep-rooted than an outward resemblance. Fundamental laws of nature have their impact both on biological organisms and technical systems. Both are born to function in the same environment. Whatever the difference between the organic and the inorganic, the identity of goals results in a similarity of structure and function. In this book, we will attempt to reveal some of the traits of this similarity.

A new element in a technocenosis can be

said to be created in the following way. In his rational activities, man never does anything "for the fun of it" or without purpose. All machines are invented, designed, and manufactured to solve a particular problem, i.e. to achieve a certain goal. These products have a purpose of existence which we shall label G_e. The first inventors of a hoe, for example, had a specific purpose: they sought to create a device for loosening soil to the depth necessary to plant seedlings. This admitted a variety of admissible forms for the working surface, the length of handle, and the choice of material for the different parts. The designers of the BelAZ lorry had a specific purpose too: they were to build a vehicle to transport large volumes of rock in opencast mines. This purpose brought into "existence" a family of superheavy trucks.

It is natural that the creator of a new object should see to it that the G_e is within reach. He also wants to achieve his goal in the most effective way. It is not easy to have a ready answer as to what this implies. It is important however that among other things, in order to achieve G_e in the most effective way, the consumption of a resource (*e.g.* energy, raw materials, or time) must be reduced to the minimum, or the job should be done with maximum probability of success and completeness or accuracy, and the cost of the change in the technocenosis which may prove necessary for the new

object to function must be minimized. All these requirements may be defined as "limitations". Although some of our readers will not agree to this meaning for the term "limitation" because it is used in a narrower sense in control theory. Its broader meaning appears justified in this book.

For a new object in a technocenosis to attain its G_e, the process for achieving it must be properly organized. It is not enough to make a hoe, you need to train human beings to put it to work. It is not enough to manufacture a BelAZ lorry, you need a driver trained to utilize it properly. In other words, the achievement of G_e requires an adequate control process. The implementation of control, in turn, demands a variety of resources designated R and information concerning the current state of the environment in which our object is operating, the state of the object itself, and the state of the controller. All this information is represented as I. Essentially, the need for control singles out two parts of an object, i.e. the one that controls and the one that is controlled. Though this division is arbitrary and in reality only proves valid at a level of discussing the object, we find it rather handy. It is the controlling part of an object that generally evokes reflections on the effectiveness of achieving G_e. To evaluate its operation the criterion of control is introduced (e.g. the achievement of a goal with the minimum

means). This is designated Q and it either may be quantitative or qualitative.

As time goes by, technocenoses become more numerous and complex and the number of interconnections and interdependences rises. This brings about an unprecedented problem of control within a technocenosis. The reason is that usually the goal of a technocenosis is not defined by an individual or a group of individuals, and the G_e of some objects which make up a technocenosis are at odds with those of others.

For example, let us assume that it is necessary to take a container with a cargo from point X to point Y and suppose that direct delivery, even by air, is for some reason impracticable. Thus, we have to ship the container by sea, then by rail and, finally, by road. Each type of transport "inhabits" its own technocenosis. In the case with sea transport, ports and their loading equipment are part of the technocenosis. The ports include wharves, loading and unloading facilities, and warehouses. Railroad transport cannot exist without sorting yards, storehouses and again loading-unloading facilities, while for road freighting the technocenosis embraces roads, repair facilities, and petrol stations. Our simple delivery operation requires a well-coordinated interaction between different technocenoses and the multitude of internal objects. If there is no warehouse-to-railroad link that can handle the container, it will

never reach its destination. If the petrol stations do not supply fuel to the lorry conveying the container and the container cannot be transferred to another lorry, the result will be the same: the cargo will never see its destination.

This coordination demands special efforts in one technocenosis or a group of interacting technocenoses*. In fact, the situation is still more complex. The point is that our man-made world does not exist independently and in fact is closely intertwined with the natural world. Man having created the technocenoses is himself their integral part, interacts with the elements of them, sets and realizes his own goals through them and, finally, coordinates things inside a single technocenosis or between several technocenoses. This involvement brings about a rapid increase in the complexity of the interconnections, restrictions, and criteria. Economic and social factors further complicate control. An impressive wave of research into the production of automated control systems has risen in the last decade, and those investigations have shown that problems of control inside and outside a technocenosis have become "hotbeds" of our technogenic civilization.

What makes the control problem which humanity is trying to solve such a stumbling

* Note that the identification of a technocenosis is in itself far from simple and is beyond the scope of this book.

block? Why haven't we seen much progress in the field so far? We shall attempt to answer these questions.

1.2. Systems That Have Never Been Designed as a Whole

There is barely no one now who has never made a long-distance call. Suppose a Mr X living in a small provincial town in the south wants to speak to his brother in a town located several thousands of miles away to the north. Once the caller X has gone through the correct procedure and his call has been connected to the person he wishes to speak to by the telephone service, neither of them cares what has made their conversation possible. What they do care about is that audibility is good and they don't have to wait long. The conversation, however, is made possible by the operation of dozens, if not hundreds, of individual devices. They make up a channel between X and Y. If the channel is semi-automatic or manual, the conversation may involve people, i.e. the telephone operators. Neither X nor Y knows the route taken by their conversation.

We are not going to focus on the control principles which allow telephones to provide a communications link between X and Y and at the same time allow for links between other subscribers. Let us concentrate on another feature of the system

which, as a part of our modern technogenic sphere, has become so familiar to so many people. In contrast to a radio receiver, a TV set, an aircraft or an automobile, the world telephone communication network was not the result of a project overseen by a single designer. The existing network has evolved from a number of simpler systems as they gradually unified and their functions became more complicated. Let us dwell on this important idea.

When the first telephone exchange appeared in the USA in 1878, it was the embryo of the future system. It certainly had a creator and a designer who also invented the method of controlling it. Thus a patching panel appeared together with the first operators, smart pretty girls who artistically manipulated the patchplugs to set up the necessary connections between the calling and called subscribers. A telephone, communication channels, and a switchboard were the vital elements of this telephone network. The local networks were quickly established in many countries. Later intercity and international channels appeared linking the individual telephone networks. This gradual structural complication of the whole network brought about new technical problems and solutions which were indispensable for the proper operation of the more sophisticated system. These improvements included repeaters, and semiautomatic and automatic switching centres.

At every stage of the world telephone system's development, inventors, scientists and designers worked to improve the network's elements, invent new commutators, raise the quality of signal transmission, etc. A hierarchy, an unmistakable indication of its structural complication appeared in the rapidly growing network. Methods for controlling the network were becoming ever more complex.

Let us have a closer look at what was happening to a communication network's control system. The creators of the first telephone networks proceeded from the idea that a subscriber should set up a connection with the operator at the exchange and tell her the number of the party to be called. The introduction of automatic exchanges did not lead to any cardinal change. Instead of telling the operator the subscriber's number, we either dial it or press keys. However, if a subscriber lives at a place not connected to automatic exchange the old method of telling the operator the city and the number is used.

Thus, the control exercised by a subscriber seeking to get in touch with someone else actually has not been changed since the time he rotated the handle of a magneto to signal to the operator. The development of the telephone networks has led only to a wider choice of potential subscribers and growth of refusals and waiting time when trying to establish contact.

No great changes have occurred in the control system at the exchange since the switching principle is still the same. There have been changes, however, at a deeper layer of control. In the telephone network, each exchange or each national company or one of these companies has its own goals and priorities. The operation of the whole network, however, requires that all the individual goals be somehow compensated and correlated to all the other individual goals. A network linking sub-networks unites all the users. The network cannot do its job if the personal interests of each user are not integrated with the interests of the others in the group. This creates a situation in which no company involved in the control of the system can maximize its own profit (by this we mean being able to satisfy the variety of demands the user makes of the network), unless it coordinates its efforts with those of the other companies. This brings in a new function of control, i.e. coordinated effort in exercising control over an object in the presence of a number of control systems (users) each having its own interests.

How do we achieve the desired coordination? There is a method implying the spread of special official information throughout the network, the organization of negotiations, and planning. However, this method is of a theoretical rather than a practical value. Negotiations and coordination would

be too time-consuming and rigid planning when user's demands are arbitrary is impossible due to dynamic and hardly predictable situations that arise in the network. What is the way out, if any? There seems to be only one. A coordinated action of control elements in the network should occur as if each element were acting "all by itself" as an autonomous and decentralized unit. To make this possible, there is a need for a control mechanism not governed by some "supreme agency" of the network and instead occurring via the exchange of information of a local character between each of the control subsystems. In a global telephone communication network, this takes the form of payments between the companies in accordance with the quality of coordinating operations in different sections of the network and the volume of the service demand requested in each section of the network.

A good understanding of these arguments is extremely important for an understanding of what follows. A system that arises as a result of evolution cannot be controlled in a centralized manner by a single control element. Thus, the global goal of a system's operation, i.e. its G_e may be achieved through the coordinated action of individual subsystems and through a conformity in the control system supervising the separate subsystems. When centralized control is imposed on such systems, the result is

inevitably a failure. The prominent Soviet scientist Lyapunov once gave a vivid example of the absurdity of centralized control in some technical systems. Suppose it occurs to someone that it is a good idea to exercise centralized control over the application of all the goods-wagons available on the Soviet Union's railroads. The extreme situation would be to create a grand goods-wagon fleet somewhere in the Urals. Transportation requests would have to come to the central control office, which would then send great numbers of empty goods-wagons to the loading sites. It is clear that such organization of control over the goods-wagon stock would be inefficient and would do nothing but harm. It does not mean, however, that a global goal and a criterion of control with regard to the goods-wagon fleet are unthinkable, though they would have to be achieved without insisting on direct control over all the resources as in the central goods-wagon fleet example. It would instead be exercised through the coordinated action of subsystems with such a system of rewards and fines that would permit the achievement of the goal with a global control criterion taken into consideration.

Such control systems are quite common. Here are some instructive examples to illustrate this.

1.3. A Few Instructive Examples

1. Everyone will have been to a market. With the exception of people who come there out of curiosity or for window-shopping, the people fall into two groups: buyers and sellers. The major operation of a market is buying and selling. As a result of this operation, a certain amount of a product passes from the seller's hands to the buyer's hands. We are not going to offer a thorough study of subtleties connected with this operation for they are beyond the scope of this book and are discussed in the literature ranging from the law and psychology to fiction. We will simplify the process and describe a single buying-selling event in terms of the following three parameters: C_1, C_2 and C. Here, C_1 is the seller's price, i.e. the minimum acceptable price for the seller; C_2 is the buyer's price, i.e. the maximum acceptable price for the buyer; and C is the transaction price that occurs in the trade if a bargain is made. In this case, $C_1 \leqslant C \leqslant C_2$. Let us assume that C_1 is not agreed on by all the sellers (though this is sometimes exactly what happens) and C_2 is not agreed on by all the buyers present. The value of C in this case will be a result of a certain process that occurs in each buyer-seller pair. However, nobody forbids a buyer to collect all the information he wants concerning the prices offered by different sellers and nobody forbids the

seller to study the prices which seem to be acceptable to different buyers. Nobody makes the buyer choose a particular seller or makes the seller choose a particular buyer. Both use changes in price as a tool to control the trade.

For clarity we will consider buyers to be the object of control and sellers to be a control system whose goal is to sell all the products available (assuming that offer and demand are balanced). This common goal is made up of all the individual goals of the sellers who want to sell all produce they have brought to the market. The common goal is of little interest to an individual seller. It is of more interest to local authorities who want to receive a percentage of the amount of produce sold in the market. This rough model suits us perfectly though in real life the percentage is obtained indirectly and depends on how the population's demand for produce is satisfied.

Now if we establish a rule for price changes during the bargaining and take into account that the seller's goal is to maximize C, make $C \geqslant C_1$, and sell as much produce as possible, it becomes clear that the goal could be achieved by replacing each seller with a machine which changes the price according to the rule beginning with some price C^* determined by the seller's initial condition on his arrival at the market (say, the value of the sale price

agreed on before the seller starts out) up to the seller's price C_1. In the simplest case, every buyer's refusal to regard the requested price as reasonable reduces the price on the basis of the rule of price changes used by the machine. Each selling machine operates in a somewhat autonomous manner. The multitude of buyers and other sellers make up an environment which tells it when to raise or depress the price. The analysis of the processes in a market shows that the decentralized control exercised by many selling machines results in a situation when C prices equalize in the market. The common goal of the whole group of sellers is also achieved for the amount of products sold out tends to be maximum. If we increase the volume of information the sellers obtain about, for instance, the average value of C at a given moment for all the trades that have occurred or provide them with all the information concerning all bargains so far made, the process of price equalization will converge at a much higher rate.

What is most significant in this example? First of all, the existence of decentralized control. The control is realized through a group of almost autonomous machines (sellers) which only obtain information about each other's actions from their environment. Paradoxical as it might sound, this method of obtaining information helps achieve all the local goals of the sellers

and ensures a certain global gain satisfying the interests of a higher-order control system without having any direct effect on the local processes occurring in buyer-seller interactions.

2. Our second example concerns a beehive and describes a situation which is well-known to an observant bee-keeper. When it grows colder and the breed is in danger, the honeybees cluster around a honey comb with the threatened progeny and make up a dense mass. The temperature in the brood-comb rises and thus the breed is saved. When it is hot and the temperature inside the hive rises above a threshold level, the bees fetch water and cover brood-cells with the breed with a thin layer of water. After this the bees start beating their wings to cause a draught thus inducing the water to evaporate and cool down the brood-cells.

In the procedure described, all bees act as individuals doing their job without any centralized control because each of them feels the critical changes of temperature in the hive and then acts to eliminate the undesirable consequences.

3. Now let us discuss a rough model of an assembly line producing cars. Each section works to produce the separate subassemblies of a future car. The subassemblies are sent to the main assembly line where they are assembled into the car. Each individual production unit may have its own main assembly line which is supplied with parts

produced in second-order workshops. As a poet chose to put it, "I make nuts to fit the bolts you are making for my nuts". This all means that efforts necessary to produce all the components for a complex object must be integrated. If a nut-producing bay smashes a monthly nut target by 200 per cent while its bolt-producing neighbour produces only 150 per cent more, it is not enough to cheer the success of the nut-makers. If you cannot increase bolt output to match the nuts, you must reduce the production of nuts to 150 per cent of the norm.

Consequently, the major problem for a car-producing plant as a whole is to ensure that there is a continuous rate of production at the main assembly line instead of maximizing the component-unit output (unless the excess units can be sold as spare parts). The managers of the auxiliary workshops may be as much autonomous and decentralized but their local goals are conditioned by the general goal of the plant.

4. Residents in large cities know about radio-equipped taxis. The driver of such a taxi can receive information about the requests coming to the dispatcher at the controller office. Now let us have a closer look at their distribution. The dispatcher has several distribution strategies to choose from. First, he may simply tell all the drivers about each request. If anyone is interested (say, he is looking for a client and

is in the vicinity of where the passenger must be picked up, or perhaps he for some reason wants to move to the destination specified in the request), he may agree to do the job and tell this to the dispatcher. Alternatively, the dispatcher may assign the request to a particular driver proceeding from his own understanding of the situation or for some personal considerations. The second method, however, proves much less efficient both for the taxi fleet and the client because the average waiting time for a taxi is bound to grow. This is even more so if the requests are not to be carried out immediately but are ordered for specified time. Here again decentralized control over taxi drivers happens to be more effective than centralized control and the dispatcher has an information-conveying function rather than a control function.

1.4. Analysis of the Examples

We believe that these examples yield certain conclusions. We shall in future chapters show the reader many more natural and man-made systems which possess the same characteristics as a telephone network, a beehive, a market, or a taxi fleet.

We have already introduced the control system parameters I, Q, and R. Now, we use the above examples to present a general classification of control systems for the narrow purpose to support the ideas we

will suggest. We do not endeavour to address the more effort-consuming task of a comprehensive classification of centralized and decentralized control systems.

Now let us reconsider temperature control in a beehive. Each bee has all the information it needs about the hive's condition. This information is the same for all the bees and it is limited to a knowledge of the temperature in the hive at a particular moment. Not only do they have similar I but also similar Q. The target of the control is to bring the temperature to a certain acceptable temperature range and the time assigned for the operation is an additional requirement. The R's are the same for all the bees. All of them can beat their wings at the same rate, raise their own temperature ten degrees above the temperature of the air, and fetch water. Thus, to adjust the temperature in a beehive, we are dealing with a control system composed of similar control subsystems each having the same I, Q, and R parameters. Besides, the local interests of the subsystems are integrated with the global control goal (indeed they coincide with it), which makes any coordination action unnecessary.

These systems are simple decentralized control systems. The control subsystems are not specialized, so if an individual control system fails to ensure control, we can obtain an adequate control in the system merely by increasing the number of sub-

systems. The motto of such systems is "If you cannot outwit, try to outnumber".

When a fire-fighting team arrives at a fire, the situation differs slightly from that in a hive. Although all the fire-fighting units have the same global goal, i.e. to extinguish fire at minimum cost and in the shortest time, and have the same initial information, their methods of work and, consequently, their individual goals are different. Some try to break the roof hiding the fire with hooks and axes, others might try to choke it with a special foam. Here we have a structural decentralization in the control system (each fire station) as well as subsystem specialization according to the methods used. This sort of systems most frequently occur in technical systems which make up technocenoses. If the fire is not attended by a chief who assigns a mission to each unit and coordinates their actions, we have a typical control system which is decentralized in the methods being used.

The assembly line of a car factory gives us an example of subsystems of a control system having no complete or similar information. The tools used in different bays and their local targets may also vary. It is the administrative subordination of all the control subsystems to the main management that brings about the centralization needed to coordinate these local targets. If there is no centralized control, cars would still be assembled if all the workers are

paid for the already assembled vehicles only. We believe the readers can imagine how this would be done.

The taxi fleet gives us an example of a system in which every taxi driver has the same information and means by which to achieve his own goals and makes different decisions due to differences in the criteria for achieving them. The dispatcher's role here is restricted to bringing these personal goals in conformity with the goal of the fleet (say, the annual plan set by upper managers). In an emergency, the dispatcher can impose his will on a driver in order to achieve the global goal of the fleet.

Finally, the situation at a market where information about bargains between sellers and buyers is lacking gives us an example of decentralized control in which the effect of a control system on the others leads to coordination between the sellers' personal goals and the equalization of prices during the trading process.

This means that in different control systems decentralization as a term implies different things. The obligatory thing, however, is the presence of individual subsystems which, although devoid of information about the decisions being made by other subsystems, have a choice about the means they must use to produce a desired effect on the object of control. Corrections to the subsystem's operations may only be induced via information obtained from the

object of control. It is not infrequent that certain subsystems either "don't know" or receive rather limited information about the existence or functions of other subsystems.

Figure 1.1. gives a classification which uses three classification parameters. An

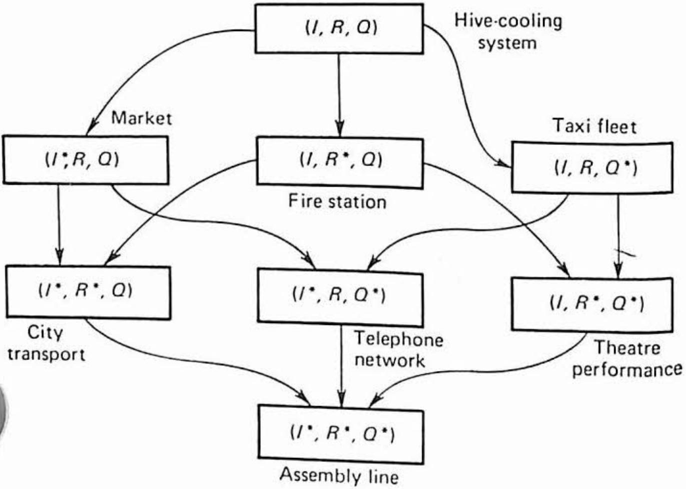

FIG. 1.1

asterisk means that in the control system various subsystems do not have the same values of a parameter. Each position in the classification is accompanied by an example of an appropriate control system. Two cases only need a special explanation, for all the others have already been analyzed. An

urban transport control system is an example of an (I^*, R^*, Q)-system. In such a system, the division into control subsystems corresponding to the types of transport (trolleybus, tram, subway) is natural. Each of these subsystems has its own way of achieving the same goal, i.e. to maximize the number of passengers carried and minimize the consumption of energy and resources. An example of an (I, R^*, Q^*)-system is the management of a theatre performance. The show's director seeks to coordinate a variety of subsystems, such as stage illumination, music, and scenery changing. All the subsystems have their own goals which are integrated into a whole by the director's plan whose purpose is to maximize aesthetic delight experienced by the theatre-goers.

In a number of cases the examples we presented were simplified to make them vivid and natural. If the reader finds the main idea of the classification in Fig. 1.1 comprehensible enough, we, the authors, will be quite content.

1.5. Why Decentralization?

A natural question arises: is it possible that the decentralization in a control process is brought about because we do not know the behaviour of an object of control well enough or because our idea of what a good control system should be is but

vague? Considerations given below are an attempt to show that in many cases decentralized control is not a defective version of its centralized or hierarchic alternative but a useful and frequently the only type of control practicable. We have had some reflections on this point already. Now let us summarize.

1. Many of the systems existing in technology, economy, and management were not created by a single designer. They were not "invented, designed, or manufactured" and instead evolved from simpler systems. The phrase we placed in quotation marks is in a way metaphorical. It is obvious that these systems are a joint effort of inventors, designers, and scientists. The point is that none of these systems was created as a whole. There have been individuals who designed some automatic telephone exchanges or local communication networks but the world has never seen a single creator of the World Telephone Network, or the World Transport System, or the World Philatelic Society. In systems of this kind, centralized control is only practicable at the level of agreements concerning the strategies of system development, standards, the restrictions to be imposed on system users, and so on. It appears, however, that operational control over such a system may only be exercised in a decentralized manner.

2. The complexity of the systems man

seeks to handle today has reached such an order of magnitude that centralized control is no longer possible due to a sweeping stream of information to be processed by the central control body and transmitted via the communication channels. As a rule, any such control proves to be so time-consuming that further control effort to manage a system dynamically would be utterly useless. The situation faced by a meteorologist trying to make a short-term weather forecast may illustrate the point. At present, ground weather bureaux and weather satellites produce so much information that processing it within the requisite time period would not be feasible. One prominent weather expert has said bittely: "Today I could tell you what the weather would be like tomorrow with one hundred per cent accuracy. But I would need a month to do the job."

3. The more complex a large-scale system is, the less reliable it is. When the number of connections is of the order of 10^{10}, as it is in the today's world telephone network, constant failures are inevitable. In other words, the system malfunctions with a probability approaching unity. Yet, the world telephone network shows no indication of collapse. The reason is that the decentralization ensures the redundancy of control necessary for the system to function properly. The dependability of such large-scale systems as the world's communication

system or a nationwide power system is provided by allowing local decisions on channel switching or the transfer of power from one part of the power system to another. Should these decisions be made in a centralized manner, the great loss of time caused by the message transmission would make the large-scale system totally unreliable and inefficient.

4. In some cases it is extremely difficult to define a control object's purpose of existence or its control criterion with the accuracy required for centralized control. Even if a definition is ventured and it happens to be successful, it would still be hardly possible to understand how to apply this knowledge to control the object. However, we may be more optimistic about having enough information concerning the control subsystems whose functions may be integrated in terms of restrictions imposed by the purpose of existence and control criterion over the object as a whole. An example is an automated system of control over a city or a region.

5. In the international and intergovernmental systems functioning in today's contradictory world, decentralization is a must even though certain intergovernmental control systems can be established by special agreement.

These considerations call attention to the decentralized control principle applied to complex systems. The first significant setups

with this type of control were created by Mikhail Tsetlin, whose contribution to this field can hardly be overestimated. He was the founder of a new branch of research known as automata collective behaviour theory. Tsetlin put forward the underlying principles for such setups and showed how to implement them. Further research in the field brought about a number of fresh ideas about the decentralized control of (I, R, Q^*)-systems and later of other types of systems.

Therefore, we shall in this book be studying various decentralized control systems (see Fig. 1.1) although we shall focus on control systems composed of subsystems belonging to one type. Here the decentralization is a result of coordination of subsystem actions through the control object, i.e. the environment in which subsystems operate. This allows the whole system to achieve its purpose as a result of each subsystem operating in order to achieve its own local goals.

This restriction on the scope of control systems we discuss is necessary to avoid a consideration of a broad class of systems which would lead to trivial results. Besides, one-type control subsystems greatly facilitate a system's arrangement and design.

Basically, we intend to dwell on such control subsystems whose operation can be described by the finite or probabilistic automaton model. Other models, however,

are also involved. We made our choice because this sort of subsystem has been the best-studied theoretically and because automata control models are the most widely used in practice.

This book is arranged as follows. The next chapter is about the model of a subsystem which is a deterministic or probabilistic automaton functioning in a random environment. We will show that a device of this type, however simple it is, can adapt its behaviour to a previously unknown environment. Chapters 3 and 4 describe the methods for promoting interaction between subsystems and for solving control problems. Numerous illustrative examples are included. In Chapter 5 we go into the problems related to homogeneous structures in which parallel and asynchronous processes occur. Such distributed decentralized control systems have many of the traits vital for today's technology and applicable when centralized control proves inefficient. In the concluding chapter we show how technological objects and their control systems have evolved. We believe this process is becoming increasingly significant in the technical progress of mankind. For systems of this kind, decentralized control appears the only possible method of control.

We have sought to write a truly popular and to a great extent descriptive text with as few analytical conclusions, proofs or

references as possible. However, we assumed that the reader would have some knowledge of the fundamentals of probability theory and mathematical analysis. Otherwise, this book would have consisted solely of the opening chapter.

Chapter 2
Is It Easy to Exist in a Contradictory World?

"Tell me where is Fancy bred,
Or in the heart, or in the head?"
Shakespeare

2.1. The Pros and Cons of Common Sense

When a fox comes back home with an abundant prey and the family has feasted, the remaining food is hidden for a "rainy day". A pit is dug with a great care, and the meat is placed at the bottom and buried. The thoughtfulness and logic of the fox's actions may deceive you into believing that these actions are generated by the animal's "intellect".

The ways of Providence, however, are inscrutable. The fox is trapped and ends up in a zoo. It is fed by the zoo attendants, and so it no longer has to waste time or effort hunting for food. However if the fox has enough food to spare, it still tries to hide it! The fox starts scratching at the concrete floor of its cage and, when the imaginary "pit" is deep enough, it "hides" the meat. Once hidden, the meat, although it is still lying on the floor, does not attract the fox's attention. It is "buried" and therefore ignored. What was useful in the fox's natural habitat does not make sense in another reality.

Very specialized actions stimulated by

a particular situation in an environment are generally called reflexes. The simpler the organism, the more rigid its reflex pattern and the more ridiculous its behaviour in a different environment. There is a vast number of reflexes and their classification is rather tenuous. Now we are going to focus only on those reflexes which help a living organism adapt to the conditions of its habitat.

Let us consider two simple examples. Zoopsychologists study the behaviour of living organisms to see how it changes in transient surroundings. They do this using mazes, the chambers and passages of which contain all sorts of stimuli which may either be pleasant or unpleasant for the maze inhabitant. By arranging these stimuli in different ways the "geography" of the habitat can be designed.

T-shaped mazes are the simplest and two of them are drawn in Fig. 2.1. Now consider the upper one. It was first used for experiments with the common worm by Yerkes, an American researcher. Initially worms were placed at the bottom of the T. This part of the maze was brightly illuminated and the worm started moving out in order to find "peace and quiet". Upon reaching the fork, it had to choose between turning left or right. It is obvious that the worm could not "know" that the left passage was fraught with much discomfort: an electric field in the way and a chamber

with an irritating solution of salt at the end. In contrast to this, the right passage led the worm into a dark humid chamber where it would feel at home.

As the experiment went on, the worm was placed in the maze many times and

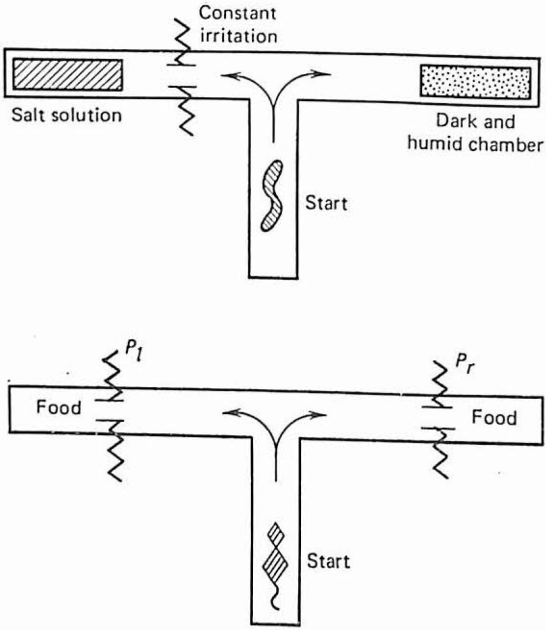

FIG. 2.1

it had to take a decision concerning the choice between these two alternatives many times. Little by little, it learned to turn right each time. In other words, a worm

having no prior information about its habitat developed a sensible behaviour by having multiple contacts with its environment.

If the investigator changed the environment by shifting the left-hand irritants to the right-hand passage and right-hand joys to the left-hand passage, that would make the clever worm's behaviour illogical. You might expect the poor thing would be totally misled and live in continual conflict with its environment. However, after several futile attempts to find comfort in the right-hand passage the worm eventually turned left. It learned the trick and adapted to the transient world.

Now we are going to consider the maze shown at the bottom of Fig. 2.1. This one was used by another zoopsychologist Thorndike, who experimented with rats. At a point where the passage breaks into two a hungry rat attracted by the smell of a bait had to choose between the same two alternatives: turn right or left. In both passages, however, the rat would suffer from unpleasant electric shocks. These irritations were applied with fixed probabilities P_r and P_l and in a given series of experiments the probabilities were kept constant. The purpose of the experiment was to determine whether the rat could gradually learn to choose the passage in which the probability of an electric shock was lower.

In experiments which followed, other

animals and different mazes were used but the main result was invariably the same, namely after a period of a more or less time-consuming training the animal would realize the difference between the two probabilities (in case of a T-shaped maze) and would make a sound decision as to which way to go to feed. When the difference in the irritation probabilities was insignificant, there was no noticeable preference in the choice of a route.

Strange as it might seem, but these facts, remarkable as they are, failed to draw the attention of mathematicians. No models were built and no discoveries made in the domain where the so far isolated sciences met. An expert in one branch of science spoke a language which was Greek to an outsider. Progress in one field was of little interest to the experts in another. In each discipline the scientists were busy studying phenomena pertaining to that particular branch. Thorndike's experiments and their synthesis and integration were separated by decades. An alliance between mathematics and zoopsychology was not yet possible in those far-off years preceding World War I. Mathematicians were blind to Thorndike's experiments while psychologists were deaf to the language of mathematics.

It took humanity fifty years to learn to look at the behaviour of worms and rats from a different angle.

2.2. A "Small Animal"

The Yerkes and Thorndike effect was first simulated and explained in a series of experiments in the 1960s by M. Tsetlin to simulate the simplest forms of behaviour. Tsetlin was an original scientist who made a large contribution to the study of behavioural simulation. Tsetlin was an inventive engineer and a brilliant mathematician, a devoted and serious student of medicine and biology and a man of many other talents. His ability to interpret facts from a variety of sciences in a precise though metaphorical manner allowed him to integrate the efforts of mathematicians, biologists, psychologists and engineers into a whole. This "invisible university" gave shape to an original scientific trend, unique for that time. Many fundamental and applied problems were solved by them, such as the creation of world's first biocurrent prosthesis. At the present moment, however, we are interested only in one aspect of their activity, i.e. the theory of collective behaviour and control.

At the base of the theory lies Tsetlin's hypothesis of simplicity. He assumed that any sufficiently complex behaviour is a combination of many simple behavioural acts. Thus, a complex behavioural process is a result of a joint realization and simple interaction of these acts. Hence, the joint effort of simple "animals" accounts for the stable existence of a whole group which

could be regarded as a "superorganism". To evoke the proper image compare it with the cells in a human body or bees in a beehive, or antes in an anthill.

Now we are coming back to Thorndike's experiment as shown in Fig. 2.2. Signals

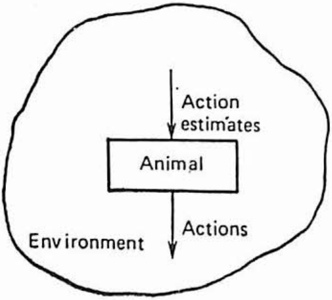

FIG. 2.2

received by a "small animal" from its environment serve as estimates of its prior actions. We can consider these estimates in a binary system: a reward (a non-fine) and a penalty (a fine). The animal can take an *action* by choosing from a certain predetermined finite set $D = \{d_1, d_2, \ldots, d_n\}$. The actions may take the values 1 and 0, the actual value being determined by the *environment*. The environments in turn differ by a manner in which the values are assigned. Consider the special case when an environment determines the values as follows. If at some moment the animal

chooses an action d_i, the environment determines whether a "penalty" (a fine) is to be paid with probability P_i or a "reward" (a non-fine) estimate is to be made with probability $1 - P_i$. If the P_i values remain constant for a certain period of time, the environment is called *stationary*. A more comprehensive definition of a stationary environment would be to specify the vector $\mathbf{E} = (P_1, P_2, \ldots, P_n)$.

Now we return to the experiment with the rat, described above. Here we are dealing with a stationary environment of the $\mathbf{E} = (P_r, P_l)$-type whose components show the probabilities of penalties (electric shocks) when a rat has to choose between the right and left passages in a T-shaped maze. The two decisions exhaust the rat's set of actions.

Tsetlin wanted to know how complex an animal must be to be able, like the rat in Thorndike's experiments, to adapt itself to a stationary environment and thus take the most expedient decision each time. However, before we attempt to give a proper answer to this question, we must specify what we mean by expedient behaviour.

Now we give our "animal" a chance to rest a little and instead we put a device for random equiprobable choice into the maze. The device disregards any fine-reward signals coming to its input and chooses either of the actions possible with equal probability $1/n$. Both passages have doors

that can be locked. Each time the rat approaches the fork it finds only one door open. The doors open with equal probability. You might even toss a coin to decide which one is to be opened. In this case the rat does not have a chance to make a proper choice, for this is done by a random equiprobable choice device.

In an infinite series of experiments with an "animal" acting like a random equiprobable choice device we will develop a total fine which is defined as the expected value and calculated by a formula which is well-known in probability theory:

$$M^* = \sum_{i=1}^{n} \frac{1}{n} P_i.$$

The value of M^* allows us to interpret expedient behaviour in the following way. We assume that the "animal's" behaviour is expedient if its total fine is less than that of a random equiprobable choice device. Thus a behaviour would be inexpedient if the total fine exceeds M^*.

Suppose that $P_r = 0.9$ and $P_l = 0.4$ in the T-shaped maze. If the rat knew these probabilities, it would invariably run along the left passage. Thorndike's experiments reveal that the rat develops this preference by experience after a period of training. If we place the rat in a situation in which the doors are opened on an equiprobable basis, the total fine will amount to $M =$

$0.5 \times 0.9 + 0.5 \times 0.4 = 0.65$. The rat's behaviour is expedient if its total fine is less than 0.65. The best behaviour possible would be to choose the left passage only and this would yield the minimum fine of $M = 0 \times 0.9 + 1 \times 0.4 = 0.4$.

The question is whether it is possible to design a device which could be as smart as our rat and behave properly in an unknown stationary environment. A remarkable achievement of the collective behaviour theory was the development of several devices that can do the job.

2.3. Reaping the Fruits of Linear Tactics

The first such machine was the *linear tactics automaton* designed by Mikhail Tsetlin. Its mechanism is shown in Fig. 2.3. Here the number of petals on the "daisy" is equal to the number of actions available to the device. For easy understanding we consider a case when this number is three. Each petal has four stable states in which the automaton may find itself. In each of these states the automaton produces a signal corresponding to the petal. A change of states is triggered by signals produced by the environment and determining the values of the actions made by the automaton. We have already pointed out that these signals are binary valued. A non-fine signal causes a change of state in the direction shown in Fig. 2.3 by the solid arrowed lines.

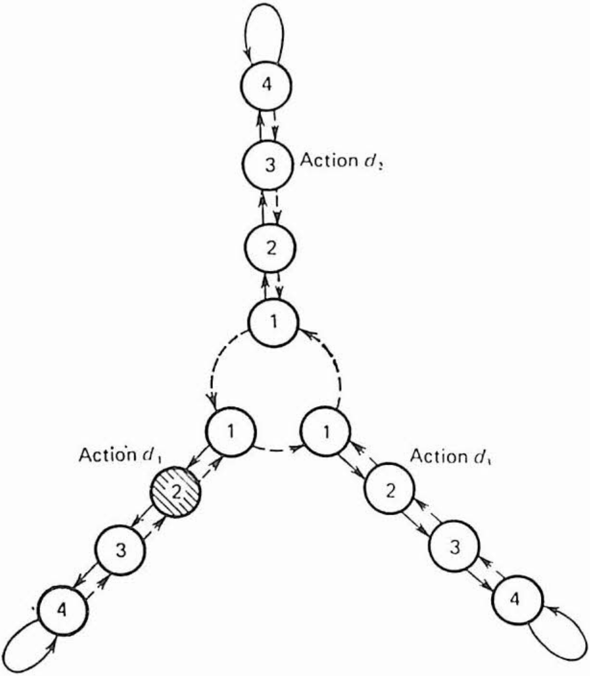

FIG. 2.3

A sequence of such signals sends the device towards the outer edge of a petal to reach the final state where it stays. If a fine signal comes to the input, the state is changed in the direction indicated by the dashed arrows. The device with a sequence of penalties moves inwards until another fine signal causes it to pass over to another

petal of the daisy; thus a change of action occurs. As can be shown in Fig. 2.3, these changes occur in a sequential fashion.

This type of automaton operates as follows. Suppose that it functions in the stationary environment characterized by the vector $\mathbf{E} = (0.9, 0.0001, 0.8)$ and that it is initially in the shaded state shown in Fig. 2.3. Now let's wait and see it work. In this state, the automaton is bound to action d_1. The environment will give the mechanism a penalty with probability 0.9 and will reward it with probability 0.1. As a result, the automaton will pass to state *1* along the same petal with probability 0.9 and to state *3* along the same petal with probability 0.1. Whatever happens, the next action of the mechanism in the environment is d_1 again. The environment inexorably responds with either a penalty or a reward signal with respective probabilities 0.9 and 0.1. From probability theory formulae applied to independent events and therefore applicable to the production of signals by the environment at each particular stage of the animal's actions, a d_1 action would be followed by two consecutive fine signals with probability $0.9 \times 0.9 = 0.81$, by two consecutive reward signals with probability $0.1 \times 0.1 = 0.01$, and by one fine and one reward signal with probability $0.9 \times 0.1 + 0.1 \times 0.9 = 0.18$. This means that after two interactions with the environment the automaton will find

itself in state *4* of the d_1 petal with probability 0.01, will stay in the shaded position with probability 0.18 and, finally, will pass to state *1* of a group corresponding to d_3 petal with probability 0.81. Increasing the number of interactions will hardly bring any qualitative change to this result. The probability of leaving the d_1 petal steadily grows with increasing interaction number, while the probability of staying in the petal declines.

What is bound to happen when the automaton passes to state *1* of the petal corresponding to the d_3 action? After completing this action the automaton is fined with probability 0.8 and passes to state *1* of the d_2 petal, or it is rewarded with probability 0.2 and passes to state *2* of the d_3 petal. As in the previous case, however, the probability of staying in this petal declines as the number of interactions grows. Consequently, the automaton will eventually have to leave the petal and pass to the d_2 petal. Here the situation is different from what we had in the previous cases. Since the probability of being fined for d_2 action is rather low, it is highly probable that the automaton will reach the last state in the petal and stay there almost forever. The probability of leaving it for the other petals is infinitesimally small. It is smaller by a factor of 10^{-15}. This means that after a certain period of training the behaviour of this automaton

will imitate a small animal and will be almost the best behaviour possible. We say "almost" because there exists an extremely small but nonzero probability that the mechanism will depart from the d_2 petal. If this happens, the automaton will have to wander through the d_1 and d_3 petals before it returns to the most advantageous d_2 petal to stay there for another long period of time. These wanderings, however, cost it an extra fine which would have been avoided had it remained in the d_2 petal all the time.

In Fig. 2.3 each petal of the daisy has four states. This number was arbitrarily chosen and a petal may contain as many or as few states as desired. We denote this number by the letter q and call it the automaton's *memory capacity*. The meaning of this variable is as follows: the larger q is, the longer it takes of the automaton to respond and the greater the number of fines is required to cause it to change its action. Intuitively we would see that the longer response time, the higher the chance of good behaviour once a correct choice for a particular environment has been made.

It should be made clear that in stationary environments the greater the memory capacity the more expedient the behaviour of the automaton will be. On the other hand, a small q would make the automaton susceptible to fine signals which would often take the device into non-beneficial petals.

Mikhail Tsetlin called this an *automaton with linear tactics*. It can be implemented with a relatively simple apparatus, with a set of shift registers, which correspond to petals, and ordinary logic gates to shift the states in these registers. Nevertheless the mechanism can cope with the complex problem of attaining an expedient behaviour in a previously unknown stationary environment. It is amazing how simple the devices capable of adaptation can be, i.e. given that adaptability at first sight would seem too hard for a simple machine.

However the behaviour of an automaton can be better than just expedient. Relying on Tsetlin's calculations we can show that, provided min P_i does not exceed 0.5, increasing q yields a series of automata with linear tactics, their memory capacity asymptotically approaching the optimum. Consequently, as q tends to infinity, $M(q, \mathbb{E}) \to \overline{M}$, where $\overline{M}$ is the minimum total fine in a particular random stationary environment. A corollary of this is that in many stationary environments Tsetlin's apparatus behaves with the greatest expediency at large q. We must admit that this sounds fantastic.

The automata with linear tactics considered above were followed by a succession of other mechanisms which were smart enough to have expedient, and often asymptotically optimal, behaviour in any stationary random environment.

2.4. Smart Machines: Reckless and Cautious

Automata with linear tactics are careful and exact. Slow but sure, they move from one state to another counting fines and rewards coming to their inputs. Other automata are more temperamental. Here is one suggested by V. Krinski. It is similar to an automaton with linear tactics in that it responds evenly to penalty signals. Reward signals, however, induce dramatic changes in behaviour, in contrast to the even and pedantic response of a linear tactics automaton, in that no matter which state a Krinski device may be in, it moves to the deepest state in the petal after a reward. Its actions are illustrated in Fig. 2.4 (the dot-and-dash lines should at present be disregarded). We could call it a "trusting" device because its optimism knows no bounds. Every non-penalty signal from the environment puts the automaton in a state of "euphoria". It might seem to be a reckless device always running into trouble. But the world of smart machines is odd. It has been proved that Krinski's trusting devices behave expediently in any stationary random environment and a succession of such devices with progressively larger memories q is an asymptotically optimal succession.

Another automaton which was suggested by G. Robbins differs from the trusting ones in that, when moving from one petal

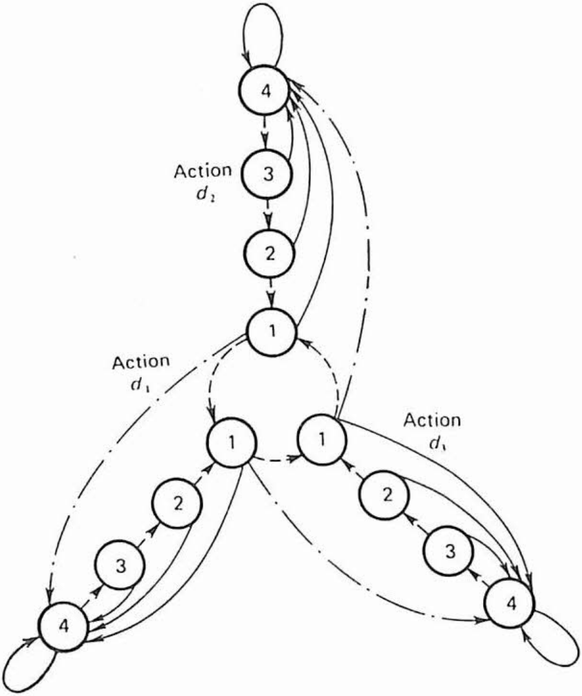

FIG. 2.4

to another, it goes directly to the final state and not to the initial state of the new petal (see the dot-and-dash lines in Fig. 2.4). Such automata too behave expediently in any stationary random environment and make up an asymptotically optimal series of automata with progressively growing q.

We have an impression that any measure that increases automaton's response times and delays in a petal's states improves its efficiency in a certain environment. For example, an inveterate angler who has found some place where the fish bite well may frequent it for a long time after he has been successively disappointed. Often his patience is generously repaid. If a change of site yields no catch the first time, he is not discouraged by the failure and will more than once return to a good old place to try his luck again. It is not until he leaves the site without any catch on a number of occasions that he is finally disappointed in it. Experience over generations of anglers shows that the average catch of an angler who returns to trusted places even after failures is always bigger than that of an angler who is easily discouraged by the first failure and is inclined to move to a more promising place.

Let us consider another type of automaton whose behaviour in any stationary environment is expedient and which can be used to build an asymptotically optimal series of automata that permit a minimum fine with any given accuracy in a particular environment. In contrast to the *deterministic* designs discussed above, this one is *probabilistic*. It has much in common with a linear tactics automaton. It responds to a non-fine signal as shown in Fig. 2.3. A fine signal, however, is not followed by

an immediate change of state. First it "tosses a coin" to decide whether to move to another state along the dashed line (see Fig. 2.3) or remain where it is. This cautious device was suggested by V. Krylov.

Another interesting investigation was to see how closely the behaviour of these devices, which were constructed using collective behaviour theory, models the behaviour revealed by Thorndike's experiments and in choice situations typical of humans. M. Alekseev, M. Zalkind, and V. Kushnarev carried out a series of experiments with people. They set up an isolated room with nothing inside but a panel with two buttons and a chair. The subject was seated on the chair and asked to wear earphones. By pressing a button the subject might hear a click, which means a reward, or not, which means a fine. The probability of a click was unknown to the subject but fixed. The subject's task was to maximize reward signals by correctly choosing between the two alternatives. The conditions are similar to those of Thorndike's experiments: a choice between two alternatives and previously unknown probabilities of reward and fine. How did people behave in this experimental situation? In the simplest case one button brought about a click with probability 1 and the other button incurred a fine with a non-zero probability. The subjects easily discovered what was up and only pressed the button which guaranteed

100% success. In more complex cases, however, their behaviour was not as simple as one might expect.

In the stationary environment specified by the vector $\mathbf{E} = (0.2,\ 0.8)$, the subjects invariably pressed first one button, then the other instead of learning to press only

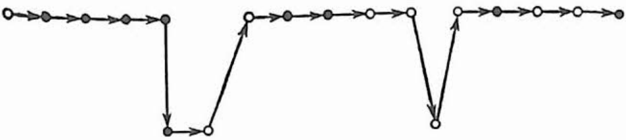

FIG. 2.5

the first button with a click probability of 0.8 even though the fine probability for the button is 0.2. One subject's response pattern is shown in Fig. 2.5. The upper chain of circles corresponds to the pressing of the first button and the lower one to the pressing of the second one. Shaded and open circles denote clicks and non-clicks respectively. As you can see, the subject mostly followed the expedient strategy of pressing the first button but this did not prevent the subject from occasionally pressing the second button. The inevitable fine which accompanied this transfer from one button (or strategy petal) to the other caused the subject to return to the optimum strategy. By comparing the behaviour of people and the operation of linear tactics automata, the investigators concluded that humans act

in the same way as automata with limited memory capacities (e.g. with q's of 1, 2 or 3). As a result, people handle the choice problem (this is especially true when the P_i probabilities are close to one another) in a worse way than linear tactics automata and the other automata we described. I. Muchnik and O. Kobrinskaya even concluded that the rats in Thorndike's experiments have much larger memories in this respect than humans. However, in environments where every action incurs penalties with probabilities which are more closely spaced, a simple unemotional automaton surpasses living beings.

2.5. How to Live in a Transient World

So far we have dealt with a stationary environment which is static and artificial, existing only in a laboratory. In real life animals exist in a habitat which is permanently changing. Survival in a transient world is far harder than adapting to a stationary environment. The laws which govern environmental parameters vary so much that merely enumerating them is extremely difficult. For this reason we will describe a dynamic environment in the following way. Consider k different environments $E_1, E_2, \ldots, E_k$ each of which is like an instantaneous photograph of a state in a dynamic environment. The projection of these photographs one after another, like

a movie film on a screen, reproduces the dynamic environment. Figure 2.6 illustrates the interaction between an automaton and such a world.

Here the commutator "switches" the animal to the environments. An automaton

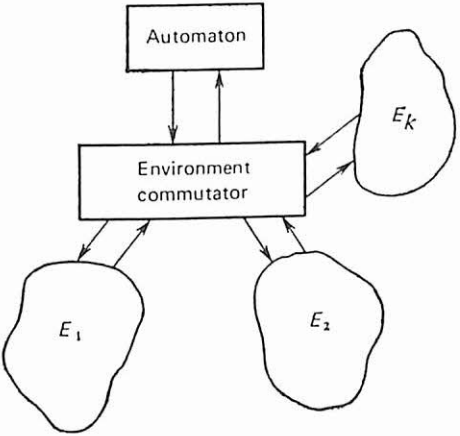

FIG. 2.6

has no *a priori* information about these environments or how the commutator operates. Adaptation now requires not only the estimation of the P_i^m values, where the superscript denotes the environment, but also a determination of the way the commutator changes between the environments.

We shall only be considering one special case in which the commutator operates. The reason we do so is that it is the best studied

case in the theory of automata collective behaviour, while the other cases are more complex and have yet to be investigated. The commutator selects the stationary environments from a matrix with k lines and k columns. The element P_{ij} is located at the intersection of the ith line and the jth column and is the probability with which the E_i environment is followed by the E_j environment after sending a signal to the input of the automaton. The element P_{ii} is the probability with which the automaton will interact with the E_i environment as it did during the previous stage. This sort of dynamic environment may be termed *alternating*. Given appropriate P_{ij} values are selected, an alternating environment can describe many dynamic environments.

What happens when our automata are put in the transient world of alternating environments? Can we still believe the postulates used to describe behaviour in stationary worlds?

Suppose we have a situation which is similar to one in a folktale. Simple Simon runs into a wedding feast and starts wailing and weeping. His behaviour is severely and instantaneously punished and he is ejected. After a while, Simple Simon meets a funeral cortege. Remembering what happened when he last met a group of people, Simon begins joking and dancing and he is punished once again. If Simple Simon were to encounter weddings and funerals in an alternating

sequence, and if he had a one-unit memory, then, as is shown in Fig. 2.7, he would have no chance of being commended. The reason is that he functions out of phase to the commutator which shifts environments from

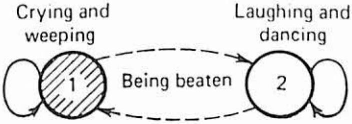

FIG. 2.7

$E_1 = (1, 0)$ to $E_2 = (0, 1)$. As a result, Simple Simon's actions are fined with a probability of one. If the environments were commutated with varying probability instead of in strict alternation, Simple Simon would occasionally be lucky and escape a beating by crying at a funeral and dancing at a wedding feast. However, he would still be beaten from time to time. The first thing suggested by our analysis of Simple Simon's shifting fortunes is that he acted as if he were an automaton with a limited memory. He had no inertia and this would seem to be helpful for automata which have to operate in random environments. But is this the correct conclusion? In a transient world the rate at which things change is so high that a response delay might be more an obstacle than an aid. The rapidly changing environments must be closely watched. Every transient world

will have its own optimal memory and this may not obey the rule "the more, the better"; the optimal memory will instead depend upon the particular rate of change. This means that we cannot even dream of designing a mechanism that will behave

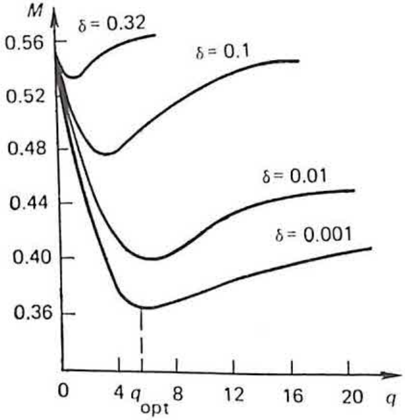

FIG. 2.8

expediently in all transient environments. This has been unequivocally proved experimentally. The results of a computer-assisted experiment are shown in Fig. 2.8. Linear tactics automata with different numbers of states in each petal were used and for simplicity it was assumed that they could choose between two actions. The alternating world was also simple. There were two stationary environments each with its own fine probabilities for the two

actions, and hence it was similar to the pair of environments in the Simon story. In the first environment, the probability of a fine for the first action was high, while for the second action it was low. In the other environment, the above probabilities were reversed. The probability of a change in the environments is denoted as δ in the figure (each curve is labelled with its parameter value). The number of memory units is plotted along the abscissa and the expected value of a total fine is plotted up the ordinate.

The experiment clearly shows that every value of δ has a memory for a linear tactics automaton which minimizes the total fine. The same result is found when other types of automata, all of which behave expediently in stationary random environments, are used in a world of alternating environments.

Consequently, the automata designs we have described so far are not best suited for dynamic worlds with fast-changing environments. The only way is to use a flexible design which can change at the same rate as the world within which it functions.

Mathematicians in the Tsetlin school have proposed several mechanisms which operate expediently in transient worlds. The best-known design is an *automaton with variable tructure* which was proposed by one of us.

Suppose that you travel to your office every day by car. You have two possible

routes and you are free to choose the one you like best. Since you leave home at the same time each day, the situation on both routes should be stationary. However, a closer examination shows that one of the routes is better than the other for it is less time-consuming thanks to less congested streets and fewer traffic lights. The trouble is that sometimes the route is completely jammed by lorries entering a nearby construction site and all you can do is regret bitterly that you chose this way on that day. Having no prior information about the lorry schedule at the trouble-making construction site, you could as well toss a coin when you leave home to decide which route to use. Time passes, however, and you learn by experience. It becomes clear that the traffic jams are very probable on Wednesday and Friday. Hence you choose the first route every day except Wednesday and Friday, on which days you go along the slower route.

This example illustrates the behaviour of an automaton with variable structure in an alternating environment. Below we present a more rigorous description of its structure and operation.

Let us go back to the linear tactics automaton in Fig. 2.3. Its structure may be specified in the form of two matrices which determine which way the states shift when a signal is received. Each matrix would contain 12 rows and 12 columns, one for

each of the automaton's states. Each row has a one to show how the change in state occurs. Since these matrices are very large, we shall consider a linear tactics automaton

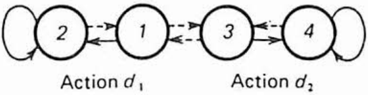

FIG. 2.9

with two states in each petal and two actions (see Fig. 2.9) instead of an automaton with four states in each petal and three actions. This automaton will have matrices of the form

$$\Pi^+ = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \Pi^- = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

These matrices define the structure of the automaton, which is deterministic. If we deal with a probabilistic automaton, like the one proposed by Krylov, then the Π^+ and Π^- matrices will contain the probabilities of the state changes rather than zeros and ones as above. For instance, if we replace the linear tactics automaton in Fig. 2.9 by Krylov's automaton, the

corresponding matrices will be

$$\Pi^+ = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

$$\Pi^- = \begin{pmatrix} 0.5 & 0 & 0.5 & 0 \\ 0.5 & 0.5 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0.5 \end{pmatrix}.$$

Unlike deterministic and probabilistic automata in which the Π^+ and Π^- matrices remain fixed whilst if operates, an automaton with a variable structure has matrices which also change. In the latter case the automaton changes its structure depending on the results of its acting, i.e. whether it incurs a penalty or a reward from the environment.

This sort of automaton starts in an "indifferent" state in which all the probabilities of changing state are the same. For the conditions illustrated in Fig. 2.9 the state-change matrices for an automaton with variable structure can initially be given as:

$$\Pi^+ = \begin{pmatrix} 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix},$$

$$\Pi^- = \begin{pmatrix} 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix}.$$

To get a clearer view of the picture we assume that the automaton's initial state is *1* and that having accomplished the action d_1 which corresponds to this state (cf. Fig. 2.9) the automaton transfers to state *4* (each change being equiprobable) according to matrix Π^+. Suppose this action is followed by a fine signal. Upon receiving this signal, the automaton which did not incur a penalty for the d_1 action decides that its $1 \to 4$ transition is a mistake. This information is recorded by reducing the probability Π^+_{14} by a value Δ. However, since the sum of the probabilities in each row of the matrix must be equal to 1, the reduction of Π^+_{14} by Δ invariably results in the growth of all the other probabilities in the row, say, by $\Delta/3$, in order to normalize the row. If $\Delta = 0.03$, then after the action matrix Π^- will remain the same while matrix Π^+ will become

$$\Pi^+ = \begin{pmatrix} 0.22 & 0.26 & 0.26 & 0.26 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix}.$$

At the next event the automaton takes the d_2 action which corresponds to state *4* and chooses the next state according to mat-

rix Π^-. This matrix is used because a penalty was the last signal produced by the environment. Suppose the automaton chooses the transition $4 \rightarrow 4$ and is again fined. Now it is time for matrix Π^- to change and for matrix Π^+ to remain the same. The fourth row of Π^- becomes (0.26 0.26 0.26 0.22). At the next event the automaton again uses the probabilistic transition in compliance with matrix Π^-. The probability values in the fourth row of the matrix will depend on the incoming signal and the next choice will be made according to either matrix Π^- (if the last signal was a fine) or matrix Π^+.

Thus matrices Π^+ and Π^- are modified by the signals from the environment. We should like to know whether these matrices tend to some stable value, e.g. to the matrices consisting of zeros and ones corresponding to a linear tactics automaton or some other automaton which behaves expediently in a stationary random environment. If so, we should be able to improve a random choice device into a mechanism which behaves expediently in static random environments. The answer to the above question depends on the way the elements in Π^+ and Π^- are changed.

It has been shown experimentally that a linear rule for changing the transition probabilities Π_{ij} in matrices Π^+ and Π^- described above does not seem to guarantee the automata proposed by either Krinski

or Robbins. The introduction of nonlinear rules for changing the elements of these matrices, however, shows that initially "flat" matrices with equal Π_{ij} tend to matrices with zeros and ones. These asymptotic matrices correspond to automata which behave favourably in stationary random environments.

This conclusion is not particularly important because in stationary random environments there is no need to waste time training variable structure automata. For these environments, well-behaved designs can be determined *a priori*. The main thing here is their behaviour in dynamic, and particularly, alternating worlds. Now we will try to see how variable structure automata fare in such environments.

Let us return to Fig. 2.8. We now know that linear tactics automata have an optimal memory capacity which depends on the rate of change of the stationary environments and which minimizes the total fine accumulated by the automata. But the memory capacity is closely related to the probability that an automaton will stay inside a particular petal and, consequently, to the probability of a particular action. A series of computer experiments with variable structure automata yielded a fundamental result. Over time, a variable structure automaton's operation in a world with alternating environments, each of which being such that a linear tactics automaton be-

haves expediently in it, asymptotically approaches the operation of a linear tactics automaton with an optimal memory capacity. In other words, a variable structure automaton finds the optimum all by itself. This is rather important because the q_{opt} value in Fig. 2.8 cannot be obtained analytically *a priori*. It must be found out as the mechanism operates, which is something a linear tactics automaton simply cannot do.

Our argument now brings us back to the ups and downs of Simple Simon. It is easy to offer many examples of alternating environments in which a linear tactics automaton would constantly incur a penalty. The moment a linear tactics automaton adapts to an environment it changes and the result is another penalty, or black eye for Simple Simon. It only needs the environments to change faster than the automaton can leave a petal and transfer to another one for a constant penalty to be incurred. If the Arctic hare were to change the colour of its fur from white to grey out of phase with the change in the seasons and if it took the hare about half a year to complete the change, it would have been extinct long ago. This situation is unthinkable for a variable structure automaton. To quote one of the first articles on variable automata, "a minimum fine is paid when yesterday's sins are rewarded and when yesterday's sins remain sins".

To sum up, we will discuss the results of an experiment involving a variable structure automaton having eight states and in a world where a linear tactics automaton

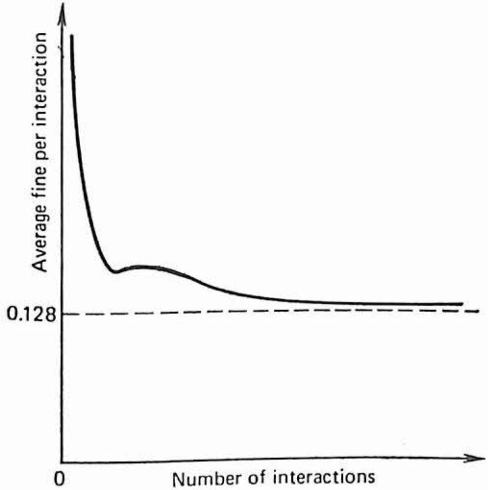

FIG. 2.10

would have an optimal memory capacity equal to two. This result is presented in Fig. 2.10. Here the number of interactions between the automaton and the environment is plotted along the abscissa and the average fine up the ordinate. The horizontal broken line corresponds to the mathematical expectation of a fine for a linear tactics automaton with a memory capacity equal to two. As can easily be seen, an initially

variable structure automaton quickly approaches the most favourable mode of a linear tactics automaton's operation and then asymptotically tends to this optimum.

The closeness of linear tactics automata to variable structure automata suggests that their designs are natural and reveals "evolutionary" links between them.

Here is another noteworthy observation. A variable structure automaton always seeks to escape a fine and enter a most favourable domain. This means it is more often rewarded than fined unless the environment is arranged so that the probability of a fine is much higher than that of a reward. In turn, this means that matrix Π^+ is more liable to changes than matrix Π^-. The automaton seems to adjust itself to function better in more favourable worlds.

We compared the behaviour of automata in stationary environments with the results of experiments in which the choices were made by people. Similar experiments were also carried out by Alekseev, Zalkind and Kushnarev for alternating environments. As the experiment proceeded, the environments were shifted without the subject's knowledge. A session comprising 75-100 button-pushing events in the environment $E_1 = (0.8,\ 0.2)$ was followed by another session of the same length in the environment $E_2 = (0.2,\ 0.8)$. The experiment gave the researchers a paradoxical result. On the average, people dealt with the adaptive

situation with an alternating environment more successfully than they handled the stationary environment. Let us have another look at Fig. 2.5. Once in a while in the stationary environment the human subjects refused to follow the better strategy because they were tempted to see what would happen if they changed their strategy. This was typical of every experimental subject. What lies behind the phenomenon can most clearly be seen when the probability of a correct choice approaches a limit. When there is a threat of a fine, refusals to stick to a more preferable strategy are less frequent. The simpler the decision, the less predictable humans become. What feature of the human mind is the cause of this contradiction? Why is the percentage of rewards in the stationary environment with $E_1 = (0.8, 0.2)$ equal to 62%, while in the alternating environment with $E_2 = (0.2, 0.8)$ it is 72%? This is only one per cent less than the result obtained in the same dynamic environment by a linear tactics automaton with an optimal memory capacity. These questions have so far not been answered. This is but another proof that frequently human behaviour is not optimal or even expedient. In this contradictory world of ours there exists a qualitative difference between a smart animal and human being.

2.6. Hungry Bats and Aerobatics

Now we will discuss whether it is possible to design a mechanism which can control the environment or adapt to it so as to maximize the mechanism's life. First we will present a general problem, then its formal description.

When a bat hunts a moth which can itself detect the bat's search signals, the moth carries out a certain pattern of actions which have been established by biologists. Experimental materials pertaining to this situation may be summed up in the following way.

The bat's vocal apparatus produces a unidirectional high frequency signal. When the signal strikes an obstacle, it is reflected. The bat receives and identifies the echo signal quickly and accurately. Thus the bat swiftly distinguishes between a static target and a moving one and between the echoes from the earth's surface and aerial objects, and it can even sort aerial objects by size into birds and mosquitoes. The bat can also determine from the echo the range and direction of potential target with great accuracy.

In turn, the moth is able to pick up the bat's locating signal, determine its source and intensity. The moth's behaviour depends on how far away the bat is and how intense the signal is. If the bat is far enough or the intensity is low, the moth carries

out a dodging maneuver. Experiments have revealed three major dodging techniques. The moth makes a U-turn and moves in the opposite direction, or climbs, or dives vertically. A closer or a more intense location signal makes the moth dart around in a chaotic flight. The latter happens because the moth's hearing apparatus is now saturated and can no longer perform its usual detection and ranging functions. The chaotic flight is a series of passive falls with folded wings, sharp turns, loops and dives. In other words, the moth follows a trajectory which makes it more difficult for the bat to predict its location from one moment to the next. We should mention that in experiments the chaotic flight strategy saved the moth's life 70% of the time.

An attempt to formalize this situation requires a minor simplification. Relying on a simplified model we are going to describe, researchers have constructed several serious models of pursuit including one which simulates the behaviour of our moth escaping from the bat.

Figure 2.11 is a graph of the state in a probabilistic automaton. Notice that for every group of states encircled by a dashed line there is a nonzero probability that it might enter a specific state (the shaded circle) in which the automaton dies. The states may be interpreted, say, in the following way: *1*—the bat is hunting and locates a moth with probability 0.3 or misses it

with probability 0.7 (for the first group of states); *2*—the bat finds the direction and range of the moth and keeps its prey in sight with probability 0.8; *3*—the bat

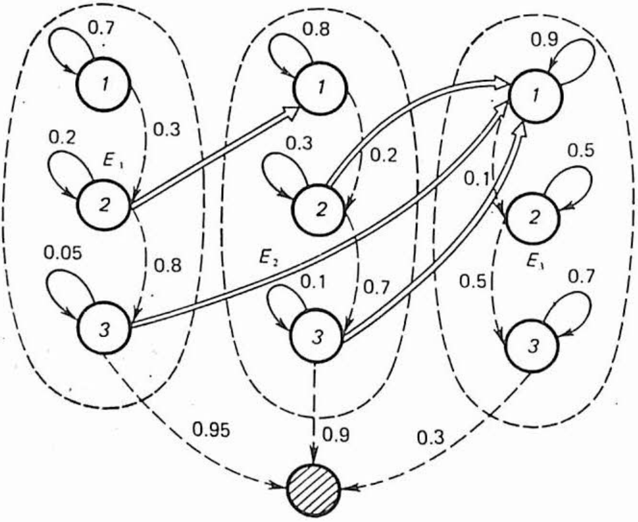

FIG. 2.11

catches and eats the moth with probability 0.95. What can the moth do in order to escape? What technique is the best? We will consider each group of the automaton's states as an environment which is specified by the strategy the moth chooses. Assume, for instance, that the three groups of states in Fig. 2.11 correspond to the following strategies: an ordinary horizontal flight (E_1),

a dive or a pitch (E_2), and a chaotic movement (E_3). The moth's actions are restricted to a change or a shift in the environments. It should be noted that the moth can only act in states *2* and *3*. In Fig. 2.11 these actions are shown by the broad arrows. In the other cases the moth carries out a neutral act, i.e. it does not change its strategy. After escaping from the bat, the moth returns to a horizontal flight, which permits it to live happily and produce offspring. These transfer actions are not shown to avoid overloading the pursuit situation.

This situation is rather simple. The actions which cause a transfer in environment allow the moth to maximize the probability of its escape from the bat. In general, however, choosing the sequence of transfers which will maximize the automaton's life expectancy is far from trivial. Let's assume, as in our example, that there are three random environments to which the automaton may shift by acting. Suppose there are three standard states and three lethal states in which the automaton dies. The first three states are numbered 1, 2 and 3, while the lethal states are denoted as 4, 5 and 6. Instead of Fig. 2.11 we will define three matrices for the automaton's change of states in the three possible environments (see Table 2.1).

Only the nonzero values of the transition probabilities π_{ij} are given in Table 2.1. If the automaton's initial state is i ($i =$

Table 2.1

Environment	States	States					
		1	2	3	4	5	6
E_1	1	0.9			0.1		
	2	0.95				0.05	
	3	0.8					0.2
	4				1		
	5					1	
	6						1
E_2	1		0.9		0.1		
	2		0.7			0.3	
	3		0.95				0.05
	4				1		
	5					1	
	6						1
E_3	1			0.9	0.1		
	2			0.92		0.08	
	3			0.7			0.3
	4				1		
	5					1	
	6						1

1, 2, 3), then the automaton's life expectancy may be calculated as

$$M_j = 1 + \sum_{j=1}^{6} \pi_{ij}(d(i))\, M_j^*.$$

Here M_j^* is the automaton's life expectancy with an initial state j provided the transition of environments by the automaton is optimal, $d(i)$ is a value of the automaton's function for the ith state, i.e. the environ-

ment to which the automaton in this state shifts from the current environment, $\pi_{ij}(d(i))$ are the transition probabilities of changes in states in the $d(i)$th environment. Apparently, the optimal transition $d^*(i)$ will occur if we have max M_j for all j's (or max min M_j).

We do not intend to take the reader through all the analytical calculations involved in the construction of the optimal transfer of environments. We should simply say that such procedure exists. It has been rigorously proved that the procedure enables a probabilistic automaton to search for an optimal environment transfer technique. We will only add for the interested reader that this procedure is a modification of Bellman's dynamic programming. In our example, an optimal shift of environments is specified by the function $d(1) = 3$, $d(2) = 3$, $d(3) = 2$ with $M_3^* = 15.47$, $M_2^* = 15.23$, $M_1^* = 13.92$. The total life expectancy of an automaton which shifts the environments is one and a half times greater than that of an automaton which lives passively. Hence, a moth which changes its flight strategies knows what it is doing.

2.7. Put Your Heads Together

We have now been introduced to a world inhabited by entities which can interact with complex environments. We must ad-

mit, however, that the models we have looked at are extremely simple descriptions of this interaction. The choice of rating (fine, reward) signals was poor and the information available for the automaton which helps it adapt to the environment was limited and the mechanisms for organizing the interactions were primitive. This was done on purpose for we sought to show that autonomous subsystems can function expediently even if they only receive meagre bits of information about the behaviour and structure of the environment. In the chapters that follow we will enrich our mechanisms by teaching them to do many more tricks. A detailed study of the abilities and evolution of such devices is, however, far beyond the scope of this book. What we are interested in is the behaviour of groups consisting of such mechanisms.

We are going to focus on the interactions between these devices and their organization into a community which can reach a common goal, i.e. mechanisms which can cause their many personal goals to agree with the common goal and distribute the resources and functions among the participants to attain the common cause.

Before we come down to the solution of these problems, let us discuss an underlying concept which provides a basis for our argument and conclusions. The reader can easily compare this concept with the models presented in Chapter 1. The major

pattern is given in Fig. 2.12. Consider a collective of k automata which interact with the environment. Each automaton functions independently and knowing nothing about its neighbours or even that they

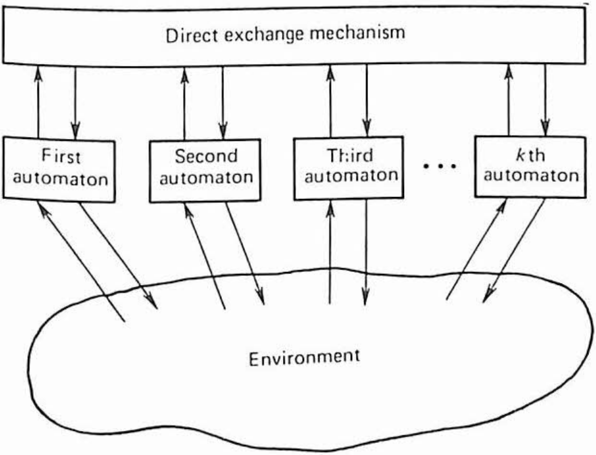

FIG. 2.12

exist. None of the automata can distinguish the other members of the collective from the environment, i.e. one automaton interacts with another as if it were part of the environment. If in one interaction the automata act locally in some way, the environment regards the actions to have a combined effect described by the set $(d_{i_1}^{(1)}, d_{i_2}^{(2)}, \ldots, d_{i_k}^{(k)})$, where the superscript is the automaton's number in the collective

and the subscript is its selected action. The environment may produce signals in response to the actions by some, all, or one of the automata. In the latter case the collective disintegrates and a study of collective behaviour is reduced to a study of k independent local behaviours. This extreme case is of no interest and we will ignore it. The other two cases are important because the environment may affect the joint effect produced by the automata.

We will also go to study collective behaviour models in which the automata "come to an agreement" with each other. In Fig. 2.12 this possibility is reflected by the existence of a mechanism by which the automata making up the collective can communicate.

Finally, it is possible to regard each of the k automata and the information exchange mechanism (if any) as subsystems of a single organism that interacts with the environment. We will use this approach in some models in later chapters.

We know that some readers may have serious objections to our model of interaction in a collective. It might seem that the rather artificial restriction imposed on the information exchange between the collective members would reduce the efficiency of the system as a whole. We would like to point out, however, that our construction is valid for the models to be discussed in this book. We shall attempt in later chap-

ters to convince the reader that our model is practical when a complete exchange of information about each automaton's actions is impossible, centralized control is absent, and the time available to choose the best action is at a minimum.

Chapter 3

"How Comes This Gentle Concord in the World?"

> "In fact, things differ greatly from what they really are."
>
> *Stanislaw Ezhi Lets*

3.1. The Sukharev Tower Pact

This chapter, which is entitled with a quotation from Shakespeare, is an attempt to show how concord may be achieved without prior agreement. It is important first to examine for a moment a situation in which prior agreement brings about disconcord.

In their satirical novel *The Little Golden Calf*, which was written in early 1930s, Ilya Ilf and Evgeni Petrov describe the activities of a group of swindlers who derive all sorts of benefits by pretending to be the sons of Lieutenant Schmidt, the leader of the 1905 revolutionary uprising of sailors in Sevastopol. In a sense, it was not easy to pretend to be a son of Lieutenant Schmidt. The idea was to extract money by applying to a bureaucrat but the danger was that he might just have been swindled by another pseudoson and thus the second impostor would find himself in a situation which, if not fatal, was more than embarrassing. What was to be done?

> "There was only one way out of the situation and that was to hold a con-

vention. Balaganov worked at the project throughout the winter. He corresponded with the competitors he knew and invited those he did not through the courtesy of Marx's "grandchildren" whom he had met in his travels. By the early spring of 1928 nearly all the known children of Lieutenant Schmidt were gathered in a Moscow apartment near the Sukharev Tower.

"The gathering was rather large. It turned out that Lieutenant Schmidt had thirty sons, ranging in age from eighteen to fifty two, and four daughters, all rather silly, rather old, and rather unprepossessing.

"In a short introductory speech Balaganov expressed the hope that the children could all come to an understanding, and would finally work out a pact, the necessity of which was dictated by life itself.

"According to Balaganov's project, the entire Union of the Republics should be divided into as many districts for exploitation as there were delegates. Each district was to be franchized to one child for ninety-nine years. No member of the fraternity would have the right to cross the region's borders or invade another's territory for the purposes of profit.

"No one objected to the new principles, with the possible exception of

Panikovsky who announced that he could manage to live without a pact. However, during the division of the territory some disgraceful scenes ensued. The contracting parties called each other names from the very first minute and addressed each other exclusively with insults.

"The arguments were caused by the distribution of the districts.

"No one wanted the university centres and no one needed the prominent cities of Moscow, Leningrad, or Kharkov.

"The sand-blown distant Eastern districts also enjoyed a very bad reputation. They were accused of being ignorant and unaware of who Lieutenant Schmidt actually was.

"... Do not take me for a fool!" Panikovsky screamed. "Let me have the central Russian plateau and I'll sign the pact!"

"... After a prolonged shouting match it was decided to distribute the districts by lot. Thirty-four pieces of paper were cut and a geographical name was inscribed for each one... . Joyous exclamations, heavy groans and oaths accompanied the drawing of the lots.

"Panikovsky's evil star exerted its influence on the outcome. He got the Volga valley. Beside himself with anger he signed the convention.

"I'll go!" he cried. "But I warn you that should I be met unkindly, I'll violate the pact! I'll cross the border!"

Every reader of *The Little Golden Calf* by Ilf and Petrov knows that eventually Panikovsky did violate the pact. But why should he have done so? Was it inevitable? Is it possible that Balaganov's pains to call the convention getting in touch with all his competitors were doomed from the start? Perhaps the trouble with Lieutenant Schmidt's children was that they were unaware of collective behaviour theory!

Let us try to formalize the situation as a game of K players. All of them are greedy and selfish and their behaviour is dominated by a desire for personal gain. Each participant has a set of alternatives which we shall call strategies and is free to choose a district to exploit by pretending to be Lieutenant Schmidt's son. The number of alternatives or strategies in the set may exceed the number of players, i.e. Lieutenant Schmidt's "children". As we have just seen, the different districts promise different profits. Each of them can be ascribed a certain number called the power of this particular strategy. In the first and simplest model we will assume that the power of a strategy, i.e. the profit that may be obtained in the district during a certain predetermined period of time is independent of the number of Lieutenant Schmidt's

children and is equally distributed between them.

This assumption means, for instance, that if one son can get 100 roubles a month from a district, then two sons wandering over the district will get 50 roubles each.

Generally speaking, the assumption is not always justified. It would be more natural to assume that the total profit in the exploitational district grows with the number of participants, the personal share of each becoming proportionately smaller. For example, when you pick mushrooms the more people there are keeping you company in your favourite spot of the wood, the fewer the mushrooms you are likely to end up with in your basket when you return home. On the other hand, the total number of mushrooms gathered by all the mushroom pickers will apparently exceed the number you could have gathered alone.

In some cases, the way the number of participants affects the achieved result is far more complex. In a moose hunt or boar hunt, for instance, each hunter's share initially grows with the number of hunters for there is a great possibility that should you venture into hunting alone, you will return home without any prey but at a certain point the catch per hunter will fall dramatically. For simplicity, however, we will rely on the first assumption.

Consider the following example. Suppose we have 10 players and the number of

strategies (or districts) is large enough, i.e. it exceeds the number of players. The power of the first district is 100 roubles a month and the power of all other districts is 40 roubles a month. Let us assume that two cunning players capture the first district and agree to split the proceeds to 50 roubles a month each, while the remaining eight players take a forty-rouble district each. In this situation, no change of district appears profitable. It is clear that none of the players will ever find it reasonable to combine with a neighbour in a forty-rouble district so long as there are vacant districts with the same potential. It does not make any sense whatever to enter the one-hundred-rouble district, which is already being exploited by two clever players because this transition will yield a decrease in power from 40 to $33^1/_3$ roubles. The same may be said about the transition from the one-hundred-rouble district with a monthly profit of 50 roubles to any of forty-rouble districts. Consequently, in this example where two players operate in a "rich" district while all the other players are in "poor" districts we have a stable situation, and no single transition promises any extra profit.

In game theory, a situation in which none of the players is interested in a change of strategy is called a *Nash equilibrium* or a *Nash point*.

It is worth noting here that, in contrast

to the actual participants, an outside observer does not care at all who settles in the bonanza district. All the situations in which two players occupy a rich district while all the rest are in poor districts are Nash points. There may be many Nash points in a game. Suppose we have one rich and twelve poor districts and ten players. There are 45 pairs of prospective owners of the rich districts, 495 ways of choosing eight poor districts and 40,320 ways in which the eight players may occupy these districts. Multiplied together, these numbers will give us a total of 898,128,000 Nash points. They have all the same total gain and equal average gains per player. The latter is termed the *value of the Nash point* or the *Nash play*.

Although it is not profitable for any player in isolation to change his strategy, it should be pointed out that the total gain and the value of the Nash play we have cited are not the maximum ones possible in this game, i.e. they may be increased. At the Nash point, the total gain is 100 roubles $+ 8 \times 40$ roubles $= 420$ roubles and the value of the Nash play is 42 roubles. If you allow a single player to operate in the rich district, the total gain would increase by 40 roubles and the average gain for all the players would be 46 roubles a month. Now look at the possibilities thus opened. At the Nash point, two players get 50 roubles each and all the rest get

40 roubles each. If the players could come to an agreement, however, a four-rouble cut in the profits of the first two players would bring about a six-rouble rise in the profit of all the rest. That's where an agreement would be useful.

On the other hand, the story shows that Balaganov's idea of drawing the lots, i.e. an arbitrary distribution of districts among Lieutenant Schmidt's pseudo children is no guarantee of stability, while a stable distribution ensures losses. Yet there must be some possibility of agreement. Or is there?

Initially let us discuss the profits to be gained by two of the Lieutenant's sons, namely, Balaganov, who contentedly wanders over his 100-rouble district, and his envious rival Panikovsky, whose evil star rewards him with a 40-rouble district. If Panikovsky violates the pact (which, in fact, is what he finally does), their profits equalize at 50 roubles because nothing on earth can induce Panikovsky to return to the Volga region. The situation in which both are engaged in their risky business in Balaganov's district is nevertheless stable, for neither can be made to go down south to the Volga region. There are at least three ways of raising their profits if only the two rogues could come to an agreement. First, Balaganov could pay Panikovsky ten roubles a month as danger money just to keep him within the unpopular Volga

district. However, Panikovsky, greedy and unscrupulous as he is, would hardly be satisfied with this sum, even if he knows that he will be unable to make much more money without an agreement. Second, Balaganov and Panikovsky could agree to a periodic change of districts which would bring them an average of 70 roubles a month. It would be natural, however, for Balaganov to distrust Panikovsky, which would make this method impracticable. Third, Panikovsky and Balaganov simply divide all the money they obtain into two equal parts. This method is termed a *common fund*. A common fund procedure or a change of districts demands a certain level of confidence between the players. Besides, the administration of a common fund itself involves additional expense. Yet, a common fund, however burdensome the expense, creates a situation with a maximum gain which is a Nash equilibrium. It is natural that what has been said above covers all the possible cases irrespective of the number of participants.

Thus we have considered what is termed an *allocation game*, and have discussed an example illustrating its equilibrium situations. This game can simulate a wide variety of circumstances. In this problem we were interested in the dependence of the players' gains on their behaviour or, in other words, the dependence of a change of strategy on the current gain.

To study the relations between the players' gains and their behaviour, it is necessary to formalize the latter by constructing a model of a player. However, what we mean by a model?

The notion is rather vague. The model of Panikovsky being thrown out of a chastened sucker's office can be a sack of sawdust, while the model of Panikovsky scheming to violate the Sukharev Tower Pact needs a much richer palette.

We have shown that what our players seek is to maximize their personal profit. Reflecting on the pros and cons of a particular strategy, they have only one criterion, which is a possibility of getting profit. Hence a model of a device optimizing its gain after a discrete set of actions could model such a player. At this point, we are bound to recollect automata which behave expediently in random environments. Such automata are devices whose "obsession" is to maximize their gains. The introduction of such automata into our game, however, requires certain changes in the game itself.

It is not hard to see that our players profit (or incur losses) in accordance with the strategy they select, while automata which stick to a particular strategy either get equal gains or suffer equal losses but with varying probabilities. We must admit that this change is not a serious obstacle. This is the more so in the case of Lieutenant Schmidt's children, who invariably de-

mand a sum of money which is obviously larger than the one they could hope for but they are satisfied with anything within the reach of local circumstances. In some cases they also suffer losses, both moral and material.

The value of the power function of a strategy is determined by the mean gain with the strategy and with a preset value of a one-time payment (either a reward or a fine). Thus, for example, if a player keeping to a certain strategy gets 200 roubles 75% of the time and loses 200 roubles 25% of the time, his average payoff with this strategy is 100 roubles. An average gain of 40 roubles occurs in a situation in which 200 roubles come 60% of the time and are lost 40% of the time.

It is easy to see that a game defined in terms of power functions for each strategy expressed in absolute gains or losses can be transformed into equivalent game in which the gains and losses have a fixed value but are determined with a probability depending on the strategy selected. Hence we are going to deal with the games with random one-time gains and losses in which the function of the players is fulfilled by automata capable of expedient behaviour in random environments. Now we have a formal model of the initial situation in which we are free to change parameters such as the characteristics of the players and the strategies and, depending on the values of

the above parameters, to investigate the course of the game as a succession of plays.

When we described automata which behave expediently in random environments, we introduced a single parameter, i.e. a memory capacity. On the one hand, it shows the automaton's structural complexity and on the other hand, its ability to average. It shows up as the lapse of time during which the automaton can take into account its gains and losses. We can assert that in our model an automata's memory capacity is an expression of a player's ability to assess the current situation or, in a sense, of their intellect.

How do game results vary with the intellectual potential of the players? It is noteworthy that an automaton's IQ is in this text reduced to its ability to average gains and losses. The players have practically no information about the game. They are ignorant of the number of other players involved, of the situation at any particular moment and even of what kind of game they are actually playing. In fact, they know nothing but their own gains and losses which allow them to choose their strategy. However primitive it might seem, it is the only way to get a clear picture of what happens during a game.

To an outside observer, there is no difference between the players. Generally speaking, an automaton whose behaviour is governed by random gains or losses will

wander randomly through a range of strategies. Therefore, we will characterize the results of the game by the mathematical expectation of an automaton's average gain, which is equivalent to the expected value of the total gain of all the automata in the game.

A thorough study of expedient automata's behaviour in this game shows that the

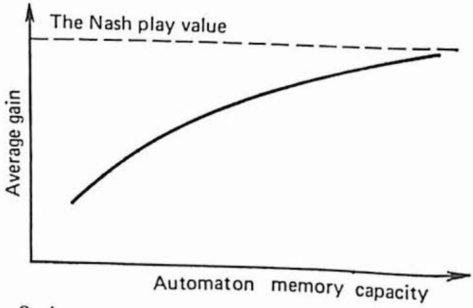

FIG. 3.1

automaton's memory capacity and the expediency of its behaviour in a stationary random environment are directly proportional to the expediency of its behaviour in the game. Hence a larger memory capacity means a larger average gain, which tends to the value of the Nash play as is shown in Fig. 3.1.

As it has been shown above, an average gain in the Nash play is different from the maximum gain which is possible in the game. We have also seen how the introduc-

tion of a common fund procedure ensures a Nash equilibrium in a play with maximum value. The Nash play with maximum value is referred to as a *Moore point* or a *Moore play*. Figure 3.2 shows how automata's aver-

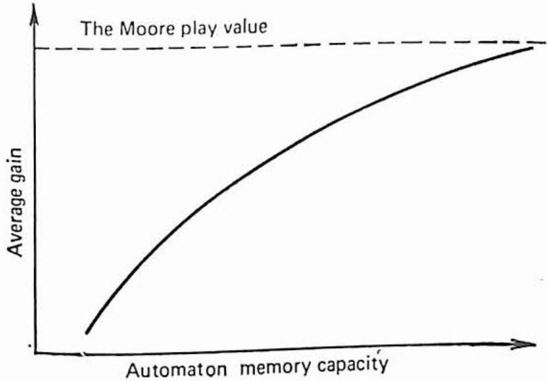

FIG. 3.2

age gain depends on their memory capacity in a common fund game. The introduction of a common fund does not seem to bring about any fundamental changes in the nature of this dependence. It is obvious that in both cases, as the memory capacity increases, the average gain rises and tends to the value of the Nash play. The difference is that in the second case the value of the Nash play is higher and is called the *value of the Moore play*. Are there any useful conclusions that can be drawn from these illustrations? Does this model give us any

interesting information? At first sight it does not seem to.

It has been shown that the Nash point requires a certain memory capacity, otherwise the players will interfere with each other and thus reduce the average gain.

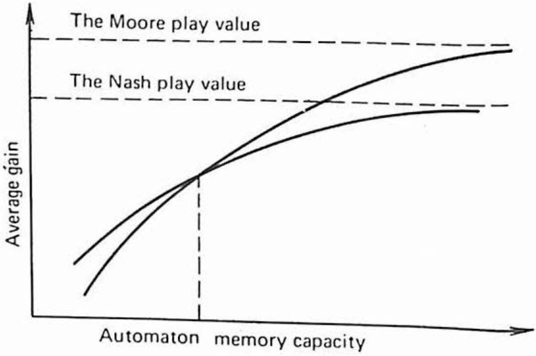

FIG. 3.3

To maximize the gain, however, it is vital to introduce a common fund, i.e. to reach agreement even if the memory capacity is large enough. A thorough analysis of the simulation results and a comparison of Fig. 3.1 and Fig. 3.2 show a seemingly obvious fact (cf. Fig. 3.3), i.e. that you can reap the benefits of a common fund procedure starting from a certain level of complexity. If the memory capacity is below this threshold, the introduction of a common fund reduces the average gain.

Tsetlin called this "the negative effect of levelling and inadequate consciousness". However, it is a matter of abilities rather than consciousness. Evidently, a common fund game demands of a player a more accurate estimate of his behaviour than a game without a common fund, where the gains and losses are connected with his behaviour by much more traceable links. A common fund procedure camouflages individual behaviour-result links.

These conclusions may be illustrated by the following example. Suppose a young girl gets a job at a chemical plant where she is to take the readings of certain instruments. The girl, like all the other workers in the plant, is given a bonus if the produce of the plant is up to the quality standards. For two solid months the girl has been reading a soul-stirring bestseller. Nothing on earth can tear her from this fascinating book; so she reads it during the work day only casting an absent-minded glance at the instruments. Luckily, nothing goes wrong and for the two months she receives her bonus. When the thriller is read to the last page, she can find no substitute and, having nothing better to do, begins taking the readings of her instruments. It so happens that, for reasons unconnected with the parameters monitored by the girl, the plant's output is flawed and no bonus is granted. In this situation, the girl is likely to come to the conclusion that

the bonus did not depend on her own actions. On the other hand, had she been punished for her own failures, it would have been much easier for her to define a proper mode of behaviour.

Now back to our model. Although simple, it allows us to arrive at a rather important conclusion: a general performance criterion is profitable only when the local decision-making apparatus is elaborate enough. If the apparatus is not complex enough, the overall aim is better achieved when each player relies on a function criterion of his own and seeks to get a larger personal gain.

Once at an international conference we were asked whether it was possible to assert from our results that capitalism is more profitable when the management system is simple and socialism is more profitable when the management system is elaborate.

Primitive as it might seem, the question is not utterly senseless. An efficient management system is indispensable for the realization of all the advantages of socialism. This is why Lenin used to say that socialism means accounting more than anything else, and why the Communist Party of the Soviet Union and the Soviet government have always concentrated on establishing a more advanced management system.

To sum up our reflections on the model of an "allocation game" we'd like to point out that this arrangement is to a certain extent reliable. When a player leaves the

game, the rest are distributed so that the vacant strategy is the one with the smallest power function. This situation may be described by the popular Russian saying "Manager fired, janitor wanted".

3.2. When Everybody is Alike

Now we will focus on a more complex model. There are two reasons why one is needed.

First, as has already been made clear, the assumption that the gains brought about by a particular strategy do not depend on the number of players keeping to it is not always valid.

Second, the simple model discussed only makes sense when there are more strategies than there are players. In fact, we assessed the results of the player's behaviour in terms of the average gain per player, in other words, the total gain of all the players. It is obvious that if the gain from a strategy is independent of the number of players who chose it, then any distribution of players among the strategies in which each strategy is chosen by at least one player ensures the maximum gain. Moreover, as the number of players grows, the possibility that the random distribution of players among the strategies will ensure the maximum gain also grows.

A model with a large number of players is again valid as soon as we make the power of a strategy a function of the number of

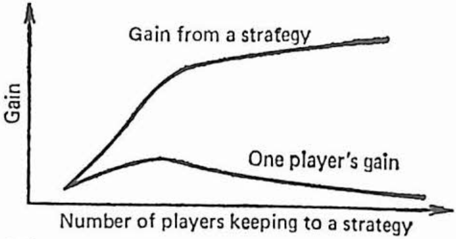

FIG. 3.4

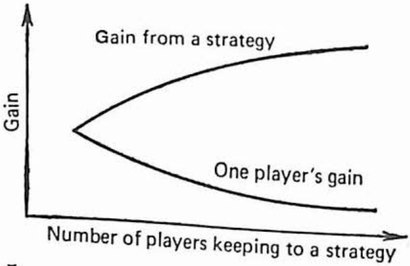

FIG. 3.5

players keeping to it. Now we consider these functions.

First of all, we may assume that the total gain in any strategy is restricted to a certain limit. From this it follows that whatever the limit, a player's personal gain should tend to zero as the number of players grows, i.e. the personal gain monotonically falls once a certain number of players choose the strategy. Typical functions of this kind are illustrated in Figs. 3.4 and 3.5.

Note that the fewer parameters specify a model, the more attractive it is. Thus we will try to construct a model ignoring

such a parameter as the number of players. To do this, we make the total gain in a strategy a function of the relative number of players keeping to this strategy instead of the overall number of players. In this case the game will be determined by the same gain functions irrespective of the number of players involved. For simplicity, we will concentrate on a case with two strategies.

Now the game is specified by two functions: the gain of the players who have chosen the first strategy as a function of their relative number among all the players and the gain of players who have chosen the second strategy as a function of their relative number. It is not hard to see that both are functions of the same variable, i.e. the relative number of players keeping to the first strategy as compared to their overall number. This automatically determines the relative number of those who have chosen the second strategy, for these are all the rest. Such a function is shown in Fig. 3.6.

We have seen how every player's gain is reduced as the number of players who have chosen the same strategy grows. On the other hand, the transition of a player, say, from the second strategy to the first one raises the gain of the players remaining with the strategy. What is the equilibrium situation in a game in which the only thing the players want is personal profit?

Figure 3.6 illustrates the situation. To the right of point a_0 the gain of every second-strategy player is larger than it would be if he were in the first strategy and a transition from the first strategy to the second one is profitable. Any single transition of the kind however reduces the relative number of first-strategy players and so shifts the gain to point a_0. To the left of a_0 the

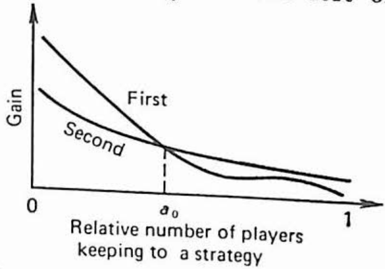

FIG. 3.6

first strategy proves more profitable and a transition from the second strategy to the first one increases the relative number of first-strategy players shifting the gain to point a_0. At point a_0 the gains in both strategies equalize. If a player chooses the second strategy at point a_0, the relative number of the first-strategy players will decrease and there will be a corresponding decrease in the gain from the second strategy, which makes this choice unprofitable. Similarly, at point a_0 it is not profitable for a player to move from the second strategy to the first one either. Hence, the distribution of

players among the strategies which corresponds to point a_0 is stable and no change in strategy is profitable to any one of the players, i.e. point a_0 in this game is the Nash point.

Consider a numerical example. Suppose there are two sites where timber is cut and 100 workers, who are free to choose where to work. At both sites, the quantity of timber stored at the site grows with the number of workers employed but their personal efficiency and, consequently, wages go down. Wages may also be depressed for different reasons including peculiarities of the management, the equipment available, the weather, etc.

Let X and Y be the numbers of workers in the first and second sites respectively. Suppose the amount of timber produced in these two sites and expressed in terms of the wages paid to the men is specified by functions

for the first site $400X-0.02X^3$
for the second site $280Y-0.4Y^2$

If the first site is operated by 80 workers, they will earn a total of 21,760 roubles, i.e. 272 roubles per person. At the second site, 20 men earn a total of 5,440 roubles, i.e. 272 roubles per person, as at the first site. The total sum earned at both sites is 27,200 roubles and none of the men would be interested in changing their place of work. It is easy to see that a single worker's

passage from the second site to the first one, for example, would cost him 3 roubles a month.

If a worker passes from the first site to the second one, his wages, like the wages of the other 20 men working at the second site will decrease by 40 kopeks. At the same time, in the first site the wages of each of the remaining 79 workers will increase by 3 roubles. Hence, total amount of money earned at both sites increases because the workers at the first site earn 237 roubles more even though the men at the second site earn 8.40 roubles less.

On the one hand, it is clear that the distribution of 80 men at the first site and 20 men at the second one is the Nash equilibrium situation, while, on the other hand, a move of men from the first site to the second one increases the total output. The total sum earned is maximized when 51 men work at the first site and the remaining 49 men work at the second one. This brings the total sum earned at the first site up to 17,748 roubles or 348 roubles per person and at the second site up to 12,740 roubles or 260 roubles per person. The total sum at both sites will be 12% more than at the Nash point and will amount to 30,488 roubles. For convenience these data are given in Table 3.1.

It is natural that when men are free to choose where to work, the optimum distribution is not stable. An 88-rouble in-

Table 3.1

	The Nash point	The Moore point
Wages at the first site	272	348
Wages at the second site	272	260
Average wages	272	304.88
Profit at the first site	21,760	17,748
Profit at the second site	5,440	12,740
Total profit	27,200	30,488

crease in a single worker's monthly wages is an incentive to go from the second site to the first one. We have already mentioned that the Nash equilibrium in a play of the maximum value may be achieved by introducing a common fund. In our example it would mean that the wages do not depend on the choice of a timber site and are determined by the total profit at both sites. In this case in the Moore play, i.e. the play of the maximum value, the equilibrium Nash wages of each worker will be 304 roubles 88 kopeks, and all this exceeds the wages at the Nash point. To ensure the equilibrium of this distribution we must pay more money to the workers at the second site than to those at the first one though all of them demonstrate the same labour efficiency. Paradoxical as it might seem, different wages for equal efficiency is prof-

itable with regard to common interests. Lower labour efficiency for some of the workers is also profitable in a situation like this.

Here we must admit that the optimal distribution of the workers among the sites may be done in a centralized manner. It is sufficient to fix the numbers of workers employed at both sites. However, this measure accompanied by an obvious inequality in wages and social problems arising as a result, may easily discourage the workers and a centralized distribution of labour resorces will be demanded. On the other hand, control over pay ensures decentralized solution of a distribution problem which will be brought about by the behaviour, cooperative or collective, of the workers themselves.

It should also be noted that our interpretation of the problem does not embrace all the situations modelled by such a game. Other illustrations may easily be found in social, economic, and technological systems.

Now we return to our examination of different behaviours in this sort of game, which we shall now term a "distribution game".

First, we would like to pass from gain functions to functions which define the probabilities of single gains or losses and make these functions dependent not upon the absolute numbers of players who have chosen to keep to a particular strategy, but

on their proportion relative to the total number of players. We covered the methodology of going from absolute gain functions to probability functions in the previous chapter. It is also not difficult to understand how we go from absolute to relative numbers. Consider the above example. Let a_1 and a_2 ($a_2 = 1 - a_1$) be the proportions of automata which at a certain play choose the first and the second strategies respectively. Let p_1 and p_2 be the probabilities of a single gain (in our example it is 400 roubles). Hence, $(1 - p_1)$ and $(1 - p_2)$ are the probabilities of a single loss of 400 roubles. The functions which specify the game will yield the mathematical expectation of a single loss with

$$p_1 = 1 - 0.25a_1^2 \text{ and } p_2 = 0.85 - 0.05a_2.$$

If $a_1 = 0.8$ and $a_2 = 0.2$, automata which choose the first strategy will win with probability 0.84 and lose with probability 0.16. If, according to our assumption, single gains and losses are equal to 400 roubles, the expected gain from the first strategy will be $(0.84 - 0.16) \times 400 = 272$ roubles which coincides with the gain at the Nash point.

Let us look at the way the gain of automata simulating players depends on their memory capacity. We shall assume that the game is played by simple automata, which maintain a strategy after a gain and change their strategy after a loss. It is

not hard to see that for such an automaton the probability of a change of action is equal to the probability of a loss in the strategy. Hence, according to the law of large numbers, if there are many automata and the probability of a loss is constant, a constant number of automata will abandon a particular strategy at each moment of time. This is true for any strategy; thus at any moment of time a certain constant number of automata will pass to a certain strategy. This may result in a dynamic equilibrium in which the number of automata leaving the strategy is equal to the number of automata adopting it. This situation is determined by a balance equation, which for our example is

$$(1-p_1)\,a_1=(1-p_2)\,a_2,$$

or

$$0.25a_1^3=(0.15+0.05a_2)\,a_2.$$

The equation has a solution for $a_1=0.63$ and $a_2=0.37$. At any moment of time 0.63 of the automata in the game will abandon the first strategy for the second one and the same number of automata will do the opposite.

The situation in which dynamic equilibrium is reached in a game played by simple automata (i.e. a game in which the probability of a change in strategy is equal to the probability of losing) is called the *Antos point*. Increasing the automata's memory

capacity decreases the probability of a change in strategy. For each level of probability for a change in strategy, however, there is a dynamic equilibrium, which, as we pointed out above, tends to the Nash point as the automata's memory capacity becomes larger. Now the average gain in a play depends on the relative position of points corresponding to the plays in a game. We shall write a_A for the proportion of automata choosing the first strategy in the Antos play, a_N for the proportion of automata choosing the first strategy in the Nash play, and a_M for the proportion of automata choosing the first strategy in a play of maximum value. Let $a_N > a_A > a_M$, which, by the way, is the case in our example when $a_N = 0.8$, $a_A = 0.63$, and $a_M = 0.51$. Then, as the memory capacity grows, the automata distribution among the strategies will move farther from a play of maximum value and toward the Nash play and the automata's average gain will decrease. In our example the automata's average gain is 299.28 roubles in the Antos play and 272 roubles in the Nash play. Figures 3.7, 3.8, and 3.9 illustrate the ways the average gain depends on the automata's memory capacity with different relative positioning of the points corresponding to the plays of the game.

Figure 3.7 features a class of games in which the most primitive automata prove to be most efficient. Of special interest is

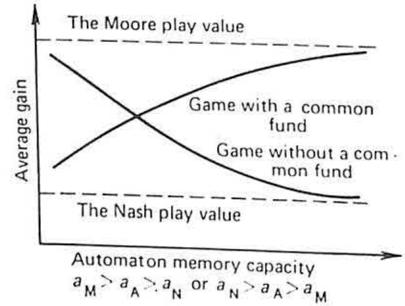

FIG. 3.7

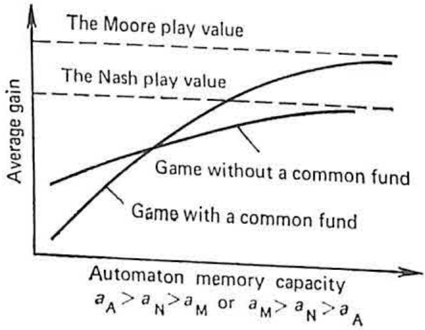

FIG. 3.8

the class of games shown in Fig. 3.9. In these games, the point of the maximum value lies between the Antos point and the Nash point. There is an intermediate and, what is most important, a finite memory capacity which ensures a play of the maximum value without a common fund.

This means that we can speak of an optimizing type of control against the background of the decentralized behaviour of the players. This control may be established by distorting the gain functions for the strategies which tend to move the play

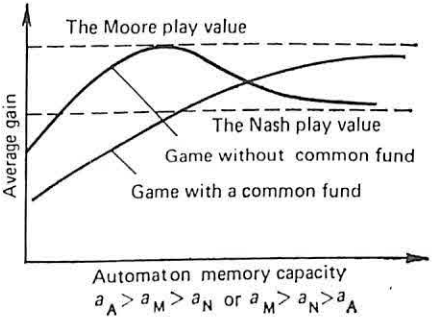

FIG. 3.9

with maximum value inside the interval between the Antos point and the Nash point. The functions may be distorted by the introduction of taxation, for instance. The reason for the success of this measure lies in the fact that the addition and subtraction of constants do not shift the play of maximum value but do shift the Antos and Nash points. Moreover, the constant may be selected so that the Antos point coincides with the play of maximum value, which means that in such a game the maximum gain will be obtained by the simplest automata.

Consider again our numerical example. If we take the 2,250 roubles earned by the men at the first tree-felling site and give this money to the workers at the second site, the Antos point will coincide with the point corresponding to the play of maximum value and the probabilities of the gain will be equal to

$$p_1 = 1 - 0.25a_1^2 - (0.05625/a_1),$$

$$p_2 = 0.85 - 0.05a_2 + (0.05625/a_2).$$

It will be obvious to a student of game theory that the above procedure is equivalent to the realization of optimal mixed strategies. Note also that this control mechanism demands that the value of the taxation constant be determined in a centralized manner.

The common fund procedure makes each player's gain independent of the particular strategy he chooses since the common fund equalizes the gains. In this case however the wages vary with the distribution of the players among strategies.

A play in which a player's gain does not depend on the strategy he has chosen but varies solely with the distribution of the players among strategies and in which the gains are equal for all those involved in the game is termed a *Goore game.*

Since the gains of all the automata in a Goore game are equal, the probabilities of a change in strategy are equal too. Hence,

in a game like this the Antos point is a play in which, whatever the payoff functions are, all automata are equally distributed among strategies. One might think that the situation in a Goore play is insensitive to changes in the automata's memory capacity because equal memories lead to equal probabilities in the change of the automata strategies. However, the procedure is governed by another mechanism and as the memory capacity becomes larger, the probability of a change in strategy decreases exponentially. The length of time an automaton remains with a strategy is inversely proportional to the probability of a change in action. The smaller the probability of losing in a certain play of the game is, the longer the automata remain with the strategy. When the automata memory capacity is sufficiently large, even a minor difference in the probabilities of losing leads to a considerable change in the probability of a change in strategy and, consequently, in the average time an automaton remains with it. A mathematical analysis of automata behaviour in a Goore play has revealed that, as memory capacity grows, automata start choosing strategies which ensure the most stable play, i.e. the ones which promise the maximum gain.

It should be pointed out, however, that to attain a maximum value play in a situation when the memory capacity is growing it is vital that the residence time in a

strategy is exponentially growing. This is true even though automata with large memory capacity achieve gains close to the optimal value due to a rapid growth of the residence time in a strategy. These considerations are useful for two reasons.

First, it becomes obvious that conclusions based on the analysis of mean values are not always valid because a similar analysis in a Goore game would tell us that automata will be invariably engaged in the Antos play.

Second, a Goore play shows up the major difficulty of optimization: it is not infrequent that the optimum situation costs so much that it is not worth it.

This may be illustrated by looking at optimization routines in computer operating systems. If a routine increases computer capacity by 5% but takes two hours to complete, the computer's actual efficiency falls by 3% even if the computer works on a 24-hour basis.

It is of special importance to bear in mind that, when dealing with on-line control problems, the time period needed to attain an optimal operating mode may be so long that by the time the transition process is over an absolutely new situation may ensue which demands that we start the whole process all over again. We have already dealt with a similar situation when we discussed alternating random environments. On-line control requires a mecha-

nism to ensure the attainment of an optimal operating mode in rapidly changing environments, i.e. at our model level which involves automata with a small memory capacity.

To sum up, we would like to point out that in a Goore play the average gain rises with an increasing memory capacity and falls with an increasing number of players. This conclusion is obvious because the more players take part in a game, the more difficult it is to analyze the nature of the relations between an individual gain and an individual behaviour, provided a common fund procedure has been introduced. Hence, this procedure is profitable for a small group of participants and it is beyond human limits to apply it to a larger group of, say, a workshop of workers.

3.3. Distribution of Limited Resources

When we speak of collective behaviour we mean the behaviour of objects acting within a framework of a system. We are obviously interested in organizing the behaviour to achieve some system goals and satisfy system performance quality criteria. At the same time, an individual object is ignorant of the general goal of the system, which brings about the need for decentralized control. What the object knows is its own local goals, local criteria, local priority functions. Control at a system level is

organized by the formation of local conditions and, probably, local interaction rules so as to ensure the achievement of the system goals by satisfying the local interests of the individual objects. Naturally we wish to know what a system object, the behaviour of which we are trying to organize, actually is.

Earlier in this chapter we considered two games: an allocation game and a distribution game (a Goore game). In both games, system efficiency depended upon the distribution of a limited number of players among strategies. To illustrate our point, we spoke about the distribution of labour resources between two places of work. In fact, the resources may be a wide range of varying objects, for example, the jobs executed in a multiprocessor system. We "personified" the types of resources and attempted to organize their collective behaviours.

Similarly, the users of resources may be looked upon as the objects which make up a system. In this case we are going to consider how to organize their behaviour so as to optimize the efficiency of the use of the resources at a system level.

The optimal distribution of resources among users only has meaning if the resources are limited. Here the term resources can mean a variety of things including money, fuel, raw materials, or equipment, etc. Each user who utilizes a certain amount

of resources achieves some effect or return. To formulate a valid problem of resource distribution it is vital that in the framework of the whole system the effects (returns) be measured by a common unit. The search for a common unit of measure is

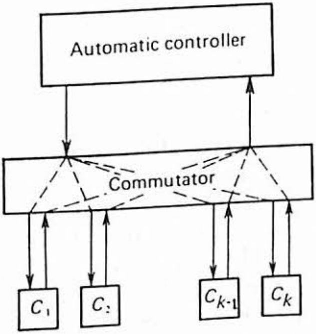

FIG. 3.10

an independent problem. In some (if not most) cases the adoption of a unit is not imposed "from above"; it comes about as the result of the joint effort of the subsystems. We will assume, however, that such a common unit does exist.

Now let us consider several examples. Suppose we have an association of k objects and an automatic controller (Fig. 3.10). The controller is connected to every object in succession through a commutator and works with each object for a period t_k with $\sum t_k = T$. The value of T is the limited

resource which is distributed among the service objects. Every object serviced produces a return equal to $\varphi_k(t_k)$. Note that all the returns obtained by different consumers are commensurable, i.e. they can be measured by a common unit, although different performance quality criteria might exist at a system level.

We may assert that the consumers are pulse servosystems which are operated by the same controller. The performance quality of each servosystem during a given period depends upon the controller's relative pulse duration and is defined, for example, by a standard deviation from the desired value. The system's behaviour is therefore given by the standard deviation of the worst servosystem. In this case, the best behaviour is min max $\varphi_k(t_k)$. It is easy to see that this criterion is satisfied when standard deviations in all the servosystems are equal. Indeed, if for a given service distribution during the period the error in one of the channels is larger than the errors in others, it would be expedient to increase the service time of this servosystem at the cost of slightly increasing the errors in the other channels. Exotic situations in which the service distribution leading to equal errors in all the servosystems is unattainable are beyond the scope of this book. It is sufficient here to assert that the errors in these servosystems are monotonically decreasing with

reductions in the relative pulse duration. The optimal resource distribution is then the solution of the system of equations

$$\varphi_k(t_k) - \lambda = 0 \quad (k = 1, \overline{k}), \quad \sum t_k = T.$$

A system performance quality criterion may also take form of a simple arithmetic sum of the returns obtained by resource users as in our example of the tree-felling sites (cf. Sec. 3.2). In the system shown in Fig. 3.10, the system effect may be determined by the total return and the behaviour of the system which approximates $\sum \varphi_k(t_k)$. It follows from nonlinear programming theory that the optimal distribution in this situation is the solution of the system of equations

$$\frac{d\varphi_k(t_k)}{dt_k} - \lambda = 0 \quad (k = 1, \overline{k}), \ \sum T_k = T,$$

where λ is the cost per unit of resources utilized. Further, we will only study two types of the resource distribution problem described above although this problem can be treated in many different ways. For example, minimizing the overall utilization of the resources for a fixed total return attained by the resource users is an important problem.

We have already pointed out that when distributing resources we are basically interested in the organization of collective behaviour under the conditions of decentralization so as to provide a solution which

satisfies the system performance quality criterion. Now we will turn to the organization of the collective behaviour of the resource users.

When organizing behaviour in this way however we cannot ignore one more participant, i.e. the owner of the resource. Can we speak of any decentralization at all when there is a central object which possesses the resource and distributes it to the users?

We are going to study decentralization of behaviour under optimal conditions; consequently, the resource-holder should not forget that the aspects of the optimization are none of his business. We will also seek to simplify information exchange reducing it, for instance, to the following procedure: the resource consumers send requests for a desired amount of resource to the centre and the centre studies the requests and divides the resource among the consumers. The simplest technique is to distribute the whole of the resource in proportion to the available requests. If x_k is the amount specified in the request of the kth consumer, the quantity of allocated resource is

$$t_k^{(R)} = T^{(R)} \frac{x_k}{\sum x_j}.$$

At this point we might naturally ask whether there are any local rules for forming requests for resources in this situation,

such that the system behaviour can be optimized according to the performance quality criterion?

Let us consider a problem with a minimax criterion. Suppose that the centre defines a value λ and informs all the resource users of this value. The users submit their requests for resources so as to make their local return equal to λ. If a consumer's return is less than λ, he may increase the request and if more, he may decrease it. If all the consumers decrease their requests, then λ is smaller than it could be, while general increase in requests indicates that λ is larger than it should be. The centre behaves accordingly: it reduces λ if the sum of the requests is smaller than the resources available and raises λ if the resources available are smaller than the sum requested. In a situation when all the returns are equal to λ and the total of the requested resources equals the resources available, the system is in a stable equilibrium. However, we have just violated the principle proclaimed above, that the centre should refrain from controlling the value of λ. It should not be overlooked that the centre must also inform the consumers about the current value of λ. On the other hand, if all the resources are distributed in proportion to the submitted requests, the resources given to the consumers are a source of information concerning the proportion of resources requested and available. This information can be

used to discover the requests at stage $(\tau + 1)$:

$$x_i(\tau+1) = x_i(\tau) - \alpha\left[\varphi_i(t_i^{(R)}(\tau)) - \frac{t_i^{(R)}(\tau)}{x_i(\tau)}\right].$$

Here the equilibrium will be the situation in which all the resources are distributed among the consumers and the returns obtained by all the consumers are alike and equal to

$$\frac{t_i^{(R)}}{x_i} = \frac{T^{(R)}}{\sum x_i}.$$

The ratio α is the consumer's "sensitivity", i.e. the length of his response time. In a sense, the ratio is akin to the automaton's memory capacity in the behaviour models we discussed above. The degree to which the optimum is attained grows as α decreases; however, in this situation it becomes more difficult for the automaton to react to changes in the environment in time.

We have already shown that if the total return resulting from the resource distribution is maximized, λ becomes the cost of a resource unit and the system performance is at a maximum when the local utility functions are the differences between the return of the resource utilization and its cost. Here the amount of resource requested

may be interpreted as the money offered to the centre to buy the resource. Having distributed the resource among the consumers in proportion with the received money, the centre thus establishes the cost of the resource unit which is the ratio of the total money sent and the number of resource units distributed. Hence, the amount of requested resource is expressed in terms of the cost of the received resource. The system performance criterion is satisfied if each consumer formulates his request so as to maximize the difference between the return on the resource utilization and the request for the resource sent to the centre. In principle, it is of no importance what algorithms or computation means the consumer uses to reach his local maximum. What is significant is that we have defined simply and unequivocally the behaviour of the centre and the local criteria which ensure a decentralized search for the extremum at a system level.

In fact a demonstration of these techniques was a major purpose of this part of the book.

3.4. What Shall We Do with Random Interactions?

In all the models considered in this chapter a player taking part in a game perceives the result of the other players' behaviour merely as the response of the environment

to his own behaviour. No automaton or player has information about the behaviour or even the presence of other participants. We made it clear above that in some situations there is no need for any additional information because automata can achieve expedient and even optimal behaviour without it. At the same time we have come across some types of behaviour which are far from encouraging, e.g. a demand for a more complex decision-making procedure (or a large memory capacity), or for lengthy periods of time required to achieve optimal behaviour. In general the term "collective behaviour" can not really be applied to such situations. Indeed we seem to be looking at models of some sort of joint effort or the behaviour of a type of "automaton gas". When we say a "collective", we generally imply a structured series of relations, the exchange of information, and cooperation between the members of the group. There is reason to hope that an account of these qualities in the sets of automata we consider in this book will, on the one hand, improve the behaviour characteristics and, on the other hand, permit a better analysis of the potential and efficiency of different cooperation techniques.

A researcher seeking to construct a model of behaviour with interaction should not overlook the fact that only simple models defined by a small number of parameters can be used to investigate the effects occur-

ring in the model and in the situations they simulate.

What sorts of interaction may be called the simplest? In our opinion, *random pair interactions* and *homogeneous interactions with limited numbers of neighbours* are the simplest type.

The main idea of a random pair interaction is that at any moment of time (at any turn during the game), the whole collective or set of automata is broken into random pairs. An information exchange occurs between each pair leading to a change in action or inner state of the automata. At the next stage the division of the automata collective into pairs is again done at random and independently of the division at the previous stage.

In an interaction with a limited number of neighbours each member of a collective is given a list of neighbours, i.e. each automaton knows the other automata with which it may interact. It is possible that such an interaction has no feed-back; an automaton receives information from its neighbours or an automaton's gain depends upon the behaviour of its neighbours while the opposite may be wrong. The homogeneity of the limited interaction arises because the sizes of the neighbourhoods are the same for all the automata. Hence, limited interaction can be defined by a homogeneous directed graph of the relations.

We start our discussion of the possibilities

of interaction from the random formation of pairs of automata.

In our study of the allocation game we have already shown that an agreement to maximize the gain may be arrived at by introducing a common fund procedure or by settling on the most profitable sites and periodically changing the sites. It is easy to achieve the same effect by a monthly procedure of drawing lots, for example. And yet, if you consider all the difficulties connected with organizing monthly conference of Lieutenant Schmidt's freedom-loving "children", you will see that this method of centralized control is not always possible. You will run into the same if not a worse trouble if you draw lots or introduce a common fund procedure by correspondence. However, you could achieve an equivalent effect to that of the common fund if the Pact were to oblige all of Lieutenant's children to exchange districts every time they chanced to meet. If such encounters were random and equiprobable, then these interactions would ensure that each member of the fraternity would stay in every district for about the same average time after a certain period of waiting, i.e. it would equalize the gains of all the participants. Thus, in order to maximize the gains it is necessary only to distribute the players initially among the most profitable strategies and then to have random pair exchanges of strategies.

It is easy to see that these conclusions may also be applied for the game of distribution. If we prescribe an initial distribution of the players among strategies and organize a random pair exchange of strategies (the first type of interaction), the initial distribution will never be violated because a pair exchange, whatever mechanism is chosen to break the players into pairs, creates a situation in which the number of players leaving a strategy is equal to the number of players entering it. On the other hand, if the process of making pairs is random and equiprobable, the players' average gains are equialized. These considerations allow us to conclude that such an interaction procedure leads to the same effect as the introduction of a common fund. It is of special interest, however, to see how the automata behaviour in an equivalent game is dependent upon the automata memory capacity.

Now we are back to the distribution game. If the automata taking part in a game have the minimum memory capacity, the interaction does not introduce any changes in their behaviour, i.e. they are at the Antos point. As the automata memory capacity grows, their behaviour tends to the behaviour in a game with a common fund while the play tends to the Moore play. An analysis and simulation of the behaviour in this situation have shown a very significant conclusion that the average gain of the

automata with this type of interaction and with any memory capacity is not smaller than the maximum gain for a particular memory capacity in an ordinary or a common fund game.

Although the first type of interaction improves on the results of automata behaviour in a game and realizes a common fund procedure without a special central agent to pool the gains and divide them among the players, it fails to improve the dynamics of the collective behaviour and convergence to the Moore point remains just as slow.

We have already shown that in the Goore play an exceptionally slow convergence arises because the dynamic equilibrium point is a point in which automata are equally distributed among strategies irrespective of the memory capacity. Moreover, in any other game and again with any memory capacity, the expected value of a change in automata distribution among the strategies moves towards the point of equilibrium distribution. We can suggest a comparatively simple random pair interaction procedure, (an interaction of the second type) which makes all the plays of the Goore game indifferent to the equilibrium with respect to the mathematical expectation of a change in the automata distribution among strategies. Once again the average gain of an automaton will depend on the time when it chooses a particular strategy.

This sort of interaction can be implemented if an automaton which seeks to change its strategy, chooses the strategy of the other automaton in the pair. If its logic does not allow the automaton to change strategy, it completely ignores its partner.

A random pair interaction of this type produces a remarkable effect. If the automata taking part in the game have a memory capacity equal to n, their average gain will be equal to the gain of automata with memory capacity 2^n in a Goore game without random pair interaction, while the convergence rate will be the same as for automata with memory capacity n in an ordinary Goore game. Note that two automata connected together and having n states make up a system with n^2 states. Since the pair has four and not two gain-loss combinations we may say that the formation of permanent coalitions of automata improves the performance quality by a power law, while random pair interactions ensure an exponential improvement.

A joint use of the both types of random pair interaction in a distribution game leads to both effects provided the memory capacity is sufficiently large. The introduction of the second type of random pair interaction, however, changes the way the simplest type of automaton behaves in the game.

We now turn to a consideration of the following situation which models the distribution game. Suppose someone wishes to

choose a holiday and there are several resorts. The holiday-maker considers a resort to be attractive or not depending on the number of people who also choose to stay at it. On the average, the more people there are, the less attractive the resort it. The less attractive it becomes the greater is the probability that the holiday-maker will go somewhere else next summer. We eat our hearts out recollecting the good old place and realizing what risk we run when trying to choose a new one. Nevertheless, we don't toss a coin or point blindfolded at the map of the world. Instead, we turn to friends for advice. The final decision comes when the wife cheerfully recounts how a Mrs McCarthy had a really good time in Brighton. The most surprising thing is that on the average and given the variety of factors, everybody is equally satisfied by their holidays. This suggests that the above procedure takes us to the Nash point, while the method for choosing a new resort looks very much like the last method for organizing random pair interaction.

Indeed, if the reader is content not to worry about cumbersome analytical calculations and is prepared to believe us, we can say that in a distribution game, the random pair interaction in which an automaton changes strategy, if it needs to, by always choosing the strategy of its partner in the pair ensures the arrival of the simplest automata at the Nash play. This is

a remarkable fact because in a game without interaction an automaton needs an infinitely large memory in order to arrive at the Nash point. A radical cut in the required memory capacity of an automaton in the game leads to a corresponding reduction in the time necessary to achieve a stationary distribution and greatly improves the characteristics of the behaviour in the case of a change in the environment. Figures 3.11, 3.12, and 3.13 (where *1* is a random pair interaction, *2* is a common fund, *3* is an ordinary game) show automata average gains as functions of their memory capacity for combined random pair interaction and the games illustrated in Figs. 3.7, 3.8, and 3.9.

In structured collectives with a fixed structure of interaction, each player's efficiency depends on what he does and what his immediate neighbours in the game do. For example, the members of a collective may be located at the nodes of a communication network or at the allocation centre for a resource. To illustrate the situation we can take a communications or a computer network in which we seek to organize the decentralized behaviour optimizing certain parameters of the system. The parameters may be efficiency, capacity, reactivity, or mean waiting time, or cost, etc. Decentralized behaviour in the solution of such problems will be studied in the next chapter. Here we are interested in some of the

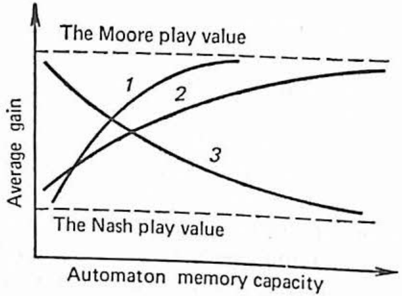

FIG. 3.11

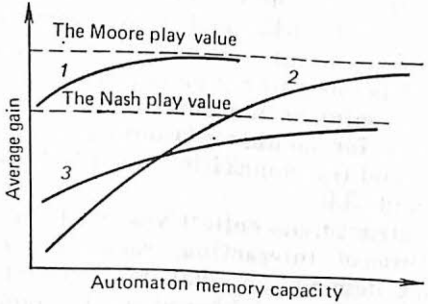

FIG. 3.12

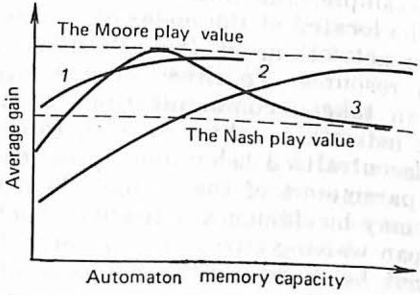

FIG. 3.13

effects brought about by the interactions which result from the structure of links in the system. A limited and homogeneous interaction is required because in the majority of real man-made networks the nodes have a limited number of interconnections and the homogeneity (as well as limitedness) must be there because it greatly simplifies the examination of the models.

The control system for an electricity or gas distribution system is an example of a controller with a network structure.

A homogeneous game with limited interaction is associated with a homogeneous graph if we can construct a function for a player's gain in terms of the strategy he has chosen and the strategies chosen by his neighbours in the game. It is natural that this may also be a function with external parameters which are uncontrolled by those taking part in the game. The homogeneity of the interaction graph means that it is sufficient to give one such function in order to specify the game.

Consider the following situation. Suppose we have a water supply network consisting of water distribution stations connected by pipelines. Each station controls the water supply to its consumers. The payoff obtained by the station is directly proportional to the total amount of water supplied to the consumers; on the other hand, a rise in the amount of water supplied may lead to a fall in the water pressure in the mains, and this

will cause losses and thus a reduction in profit. Note that the payoff is not only a function of the behaviour of the station itself, it also involves the amount of water extracted by the station's nearest neighbours. A similar situation is typical of irrigation systems.

We admit that this interpretation of a model of a game with limited interaction only gives a vague idea of what really happens in such systems, and yet we plead for the reader's indulgence. Generally speaking, gain functions can take into account all the intricacies of the performance quality criteria of the station, for example, the refusal to switch off a pump and so save power. However, we are only interested in how the gain of each participant depends on his own behaviour and that of his immediate neighbours.

In such a game there are Nash equilibrium situations at which none of those taking part in the game can benefit from a single change in their own individual behaviours. Similarly, the payoff of the whole system at the Nash point may be much less than the maximum gain. To construct a maximum value game, you could introduce a common fund. However, it is perhaps easy to see that in large systems pooling gains practically denies the players rapid information about the system's reaction to their own behaviour. A growth in the size of the system is accompanied by more complex

centralized control and adds arguments in favour of decentralized systems.

The numerical example we are going to discuss is simple: suppose two of the players in a game both have two neighbours, i.e. the graph of their interaction may be constructed as a circle (Fig. 3.14). The gain

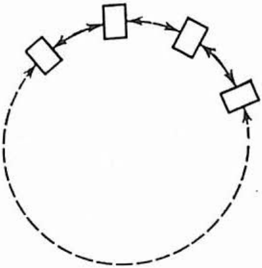

FIG. 3.14

obtained by each player depends on the strategy of his neighbours. Each player may follow one of the two strategies denoted A and B. The values of the automaton's gain depend on the strategies of his two immediate neighbours. Thus

Situation	AAA	BAA	ABA	BBA	AAB	BAB	ABB	BBB
Gain	−2	2	0	10	10	0	2	−2

From this it follows that a change in strategy is profitable to an average player if he is in situations AAA, BAA, ABB, and BBB and is not profitable in any other situation. Consider the situation BBAA, in which a change in action appears profit-

able to the third player. The result is the configuration BBBA in which a change in strategy is beneficial to the second player. The Nash equilibrium situation here is play ABABAB ... AB. The average gain in the Nash play for this game is equal to 0. On the other hand, play AABBAABB ... AABB ensures an average gain of six but, as we have seen, this is not a stable situation.

We'd like to draw the reader's attention to the following fact: if any one of the players changes his strategy, it changes his own gain and the gains of his immediate neighbours but it has no effect whatever on all other players. Hence, if we organize common funds for neighbours in the game, a single player's strategy change that decreases the total gain in the neighbourhood and, consequently, the gain for the whole chain, proves to be unprofitable. In this case the maximum value play becomes a Nash equilibrium play, i.e. a Moore play. To illustrate this we turn back to our example. The first column in Table 3.2 contains

Table 3.2

Fragment of the Moore play	Gain	Fragment of new play	Gain
AABBA	22/3	AAABA	8/3
ABBAA	14/3	ABAAA	0
BBAAB	22/3	BBBAB	8/3
BAABB	14/3	BABBB	0

fragments of the maximum value play, the second column contains the gain of an average player in the fragment with a local common fund, the third column contains the fragment resulting from a strategy change by an average player, and the fourth column contains the gain of an average player in the fragment if there is a local common fund in the new situation.

It can be seen from Table 3.2 that a change in strategy is not profitable to any of those taking part in the game when a local common fund is introduced into a maximum value play.

A local common fund, in fact, distributes the gains from each node among all the adjacent nodes. While on the one hand, this procedure does not require complicated management, on the other hand, due to the limited number of neighbours it only mildly conceals the dependence of the gain on personal performance quality. We would like to emphasize again that this effect occurs for any size of network.

3.5. He Thought I Thought He ...

The English poet Coventry Patmore wrote the verse:

> "I saw you take his kiss!" "Tis true."
> "O, modesty!" "Twas strictly kept:
> He thought me asleep; at least I knew
> He thought I thought he thought I
> slept."

The verse is a vivid example of the human brain's ability to see things from another person's viewpoint by putting oneself in another's place. Reflexive considerations are recursive, i.e. they look as if they fit inside each other like a matreshka (the Russian doll which has a series of successively smaller dolls that fit inside each other). For instance, you may reflect on how somebody else reflects on you or how he models your reflections of him. Coventry Patmore in the poem gives us a brilliant example of a recursive character of reflexive considerations.

What can we do with reflections like this? We need them when we make a choice whose success or failure depends not only upon our own decision but also upon the decisions of the other people with whom we are associated. An example of this is an allocation game in which each player's gain is determined both by his own strategy and by the strategies of the other members of the group. For this reason mechanisms which simulate reflexive considerations may be very helpful in understanding collective behaviour better. We shall attempt to show that this is so.

We introduce, using induction, the important notion of a *level of reflex*. Let's assert that an individual or automaton has a zero reflex level if his choice of strategy does not take into account the presence of other members of the group. Thus the

choice depends only upon the information which is presented to the input of the decision-making unit by the environment. An individual or an automaton has a first level of reflex if he believes that all the other members of the group have a zero reflex level which enables him to take all decision-making upon himself. Note that the presence of the first reflex level comes about for a want of information about at least some of the other members of the chain and the signals arriving at their inputs from the environment. The subsequent levels of reflex are defined in a similar fashion. An individual or an automaton has the kth level of reflex if he believes that all the other members of the chain he knows have a reflex level of $k - 1$ and he decides accordingly.

This definition of reflex levels is connected with the extent to which the decision-making system is informed about the signals received by the other systems. In case of human beings, reflexive considerations mostly rely on the knowledge which is stored in his "model of the world". This knowledge contains information about the behaviour of the members of his society, the extent of human abilities in a particular situation, the accepted conventions, restrictions, etc. Even in this simplified form reflexive considerations can be applied to a number of collective behaviour models.

Consider the following problem. A well

is sunk by a group of villagers. In each villager's garden there is a pump which conveys water from the well to a round water collector surrounding all the gardens (Fig. 3.15). The throughput of the pumps

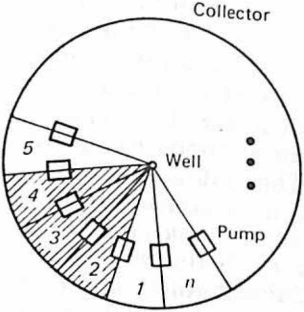

FIG. 3.15

is such that sufficient head of water can be created in the collector to water three adjacent gardens if only two pumps are working. In other words, if pumps *2* and *3* are operating, garden *4* may be watered too. Each villager has the personal goal of watering the plants in his garden. There is, however, a village trust for all n gardens which has a goal of its own, i.e. to save power. From the trust's point of view it is a waste of power to have n working pumps. Here the best situation is when the number of working pumps is $n/2$ if n is even or

$(n + 1)/2$ if it is odd. The irrigation is ensured by having only the odd or the even pumps working.

Surely, the trust could impose centralized day-to-day control over the water resources and power but the villagers assert that it has no right to do so. In an attempt to handle the situation, the trust seeks to save power by imposing fines on villagers for wasting power.

Before discussing the fines, we admit that the situation is artificial. The only thing we want it for is to show you a model of the usefulness of a reflexive way of thinking.

Let us return to the model. Suppose we have a chain of n automata (for simplicity we assume that n is even). Each automaton may be in one of two states: on or off. Again for simplicity we will label these states 1 and 0 respectively. Each automaton has information about its own state and the states of its two neighbours. Each automaton may carry out two actions. These actions are merely sending signals about the automaton's state. Each turn the automata's inputs receive either a reward or a fine signal. If it gets a reward signal, the automaton remains in the same state, while a fine signal makes it pass into the other state. The automaton's interaction in the ring with the environment (the village trust) is given in Table 3.3.

If an automaton chooses which action to

Table 3.3

States			Probability of a fine
Automaton's own state	Left-hand neighbour's state	Right-hand neighbour's state	
0	0	0	1
0	0	1	0.5
0	1	0	0.5
0	1	1	0
1	0	0	0
1	0	1	0.5
1	1	0	0.5
1	1	1	1

carry out next on the basis of the table alone, it has a zero reflex level. If all the automata in this collection possess zero reflex levels, then the trust may find itself in a situation when its goal can't be reached. If, for instance, all the automata are initially operating, all of them (according to the last line of the table) will be fined and pass into an "off" state. However, the whole collective will then get a fine signal again, and so all the automata will turn on and the cycle is closed. The pumps in all the gardens will be either on or off all at the same time and thus the trust's goal will never be reached.

Now let us use different reflex levels. Suppose, for example, that one automaton has the first reflex level. Then it will choose its strategy by analyzing the strategy changes

its neighbours will make (for this it must have information about its neighbours' neighbours) given that they have a zero level reflex, i.e. their choice is defined by Table 3.3, and then it will change state accordingly. In this case, the probability of a fine is no longer determined by the environment but by the automaton itself. In other words, an automaton with first

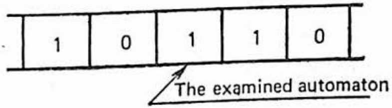

FIG. 3.16

reflex level must have information about the states of its immediate neighbours and their adjacent automata and know what is in the right column of Table 3.3. Initially the automaton studies the situation from the viewpoint of its left-hand neighbour. It is only by using the additional information that it can keep to a proper reflexive way of thinking as shown in Fig. 3.16. It follows from the table, which was compiled for the actions of an automaton with a zero reflex level, that the left-hand neighbour will not be fined at all and hence it will remain in the 0 state. The right-hand neighbour, by contrast, will change its state with probability 0.5 and so will remain in its present state also with probability 0.5. What should be done in a

situation like this? If the right-hand neighbour changes state, then the middle automaton would be in a favourable position (i.e. not fined with certainty) if it remains as it is. If the right-hand neighbour does not change its state, the middle automaton may be fined with probability 0.5. If the

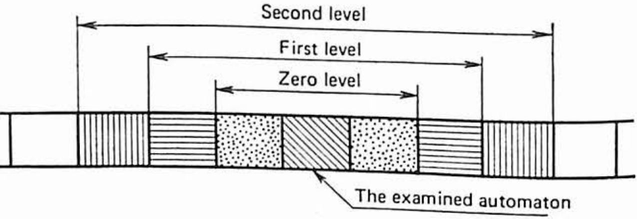

FIG. 3.17

middle automaton changes state, it will either be fined with probability 1 (if the right-hand neighbour changes state) or with probability 0.5 (if the right-hand neighbour remains in its present state). An automaton with first reflex level should therefore remain in its current state.

If our automaton has the second reflex level, then, in accordance with our definition, it would regard its neighbours as automata with first reflex level and, when analyzing their future action, it would have to account both for its own neighbours and their neighbours and also for the neighbours of its neighbours' neighbours. Figure 3.17 shows how the number of

automata whose current states must be taken into account for each reflex level grows.

Note that even though an automaton has a certain level of reflex, this does not at all mean that it can properly predict the actions of all the automata involved in the analysis. Sometimes it may make a mistake. An automaton with the first reflex level may proceed from the assumption that its adjacent automata have a zero level and that they will act accordingly. It is probable, however, that the automaton will find itself in the company of automata with levels which are higher than zero. Thus misled, our automaton will face unpredicted contingencies.

You may ask whether it is possible for there to be distributions of reflex levels among the automata around the chain such that in time the village trust can expect the whole group to arrive at a favourable state, i.e. a sequence of alternating states 1, 0, 1, 0, ... Computer-assisted simulation of this problem has shown that there are such sequences only when the reflex levels are distributed according to certain patterns. For example, the global optimum is achieved when automata with the first and the zero levels of reflex alternate throughout the chain. There is a global optimum, however, if the distribution is not strictly alternating.

At the end of Sec. 3.4 we considered a model which is similar to the one described

above. There the equilibrium situation was a play ABAB ... AB or, in the new notation, 1010 ... 10. In the pump problem this play suits us perfectly, but in the model studied in Sec. 3.4 the players arrived at the point due to the system of gains shown in Table 3.2. This table is non-existent in the pump problem. There the automata chain does not have an equilibrium point.

An equilibrium point arises due to inhomogeneities in the automata chain brought about by different levels of reflex. It is this inhomogeneity that allows us to find an optimization solution which we could not do for a homogeneous group, each individual acting on its own.

3.6. Optimists and Pessimists in the World of Automata

Now we turn to another way of introducing inhomogeneity into a collective of automata which are operating to achieve a certain aim. As always, we will start with an illustration.

Suppose a man has decided to marry. Since he regards the change in his marital status as a major turning point in his life, he believes he won't choose a wife until he has certain information about her. Suppose that he, of all things, wants to know whether she has a flat of her own and whether she is a good cook. Please do not object to

this business-like approach. We do not recommend it and would perhaps be shocked by someone who ignores the more romantic sides of love. However, in order to find a good example it is sometimes necessary to be blind to personal likes and dislikes. The way the bachelor is informed of these two aspects will be expressed in the following way. If a prospective wife has a flat, then $X_1 = 1$, if not, $X_1 = 0$. If the bachelor does not have any relevant information about whether a candidate has a flat, then $X_1 = 0.5$. Similarly, we assert that her ability to cook an eatable dinner is expressed as $X_2 = 1$, and if she is no cook, then $X_2 = 0$, while for a lack of information $X_2 = 0.5$. The variable Y will express the bachelor's decision. If he thinks she is marriageable, then $Y = 1$, if he refuses to marry, $Y = 0$ and, finally, $Y = 0.5$ if he is at a loss and still thinking what to do.

Table 3.4 illustrates this situation. It specifies five ternary logic functions which depend upon two arguments, X_1 and X_2. The simplest is function Y_1. As can be seen from the table, $Y_1 = \min(X_1, X_2)$. In logic, such a function is termed a conjunction. If the bachelor relies on this function, he agrees to marry only if both demands are satisfied, i.e. if a prospective wife has a flat of her own and her dinners are tasty. A failure to satisfy either of these two conditions causes him to refuse to marry the girl. If there is some uncertainty in a sit-

Table 3.4

X_1	X_2	Y_1	Y_2	Y_3	Y_4	Y_5
0	0	0	0	0	0	0
0	0.5	0	0	0	0	0
0	1	0	0	0	0	0
0.5	0	0	0	0	0	0
0.5	0.5	0.5	0	0	0.5	1
0.5	1	0.5	0	0.5	1	1
1	0	0	0	0	0	0
1	0.5	0.5	0	0.5	1	1
1	1	1	1	1	1	1

uation when all other demands are satisfied or if the situation is completely vague ($X_1 = 0.5$, $X_2 = 0.5$), the bachelor hesitates and says neither yes nor no. It is possible that he is waiting for more information. Such behaviour may be called objective or impassive.

The other functions in the table describe other methods of making a decision. Functions Y_2 and Y_3 reflect a pessimistic viewpoint. By accepting them the bachelor believes that he lives in a transient world in which the best thing to do is to be constantly on the alert. To him, the lack of information is a bad sign comparable to negative information. Obviously, this sort of person is a pessimist. This is even more true if he is dominated by function Y_2; his pessimism is then at an extreme. Any ambiguity

causes the person to refuse to store further information. At this point, the love affair comes to an end. Why the Y_3 function is pessimistic is less obvious. Here our anxious bachelor is inclined to kiss his sweetheart good-bye if the situation is completely vague. Whenever there is a ray of hope and she is good at one thing, he continues to collect information.

The other two functions represent the opposite world outlook. This is the view of an optimist who hopes for the best. An optimist whose choice is governed by function Y_5 is an example of an incorrigible if not "unbridled" optimism. He is prepared to wait a little in hope to see all 0.5 estimates grow into 1. Function Y_4 is the behaviour of a more cautious man who does not expect more than one 0.5 estimate to grow into a 1.

Like the levels of reflex, we can introduce the levels of optimism-pessimism. We can define the objective bachelor as having a level of zero. A person who changes m or more 0.5 estimates for 0 has a level of pessimism $n - m + 1$, where n is the number of conditions taken into account, while a man who changes m and fewer 0.5 estimates for 1 has a level of optimism equal to m. If there are two arguments as in our table, there are two levels of pessimism and optimism, i.e. 1 and 2. The number of possible levels is linearly proportional to the number of arguments n.

When deciding whether to marry a girl or not, a bachelor may be guided by considerations other than the ones described above. It is probable that the bachelor will not be as categorical as the one above and may be content with either a good flat or a good dinner regarding the concurrence of both as a happy coincidence and a godsend. The decision-making procedure of the bachelor may be represented by Table 3.5.

Table 3.5

x_1	x_2	z_1	z_2	z_3	z_4	z_5
0	0	0	0	0	0	0
0	0.5	0.5	0	0.5	1	1
0	1	1	1	1	1	1
0.5	0	0.5	0	0.5	1	1
0.5	0.5	0.5	0	0.5	0.5	1
0.5	1	1	1	1	1	1
1	0	1	1	1	1	1
1	0.5	1	1	1	1	1
1	1	1	1	1	1	1

In logic, the function $Z_1 = \max (X_1, X_2)$ is called a disjunction. It defines the choice to be made by a person who is content if only one of his requirements is fulfilled. The functions Z_2 and Z_3 reflect a pessimistic viewpoint, while Z_4 and Z_5 imply that of an optimist. As in a conjunctive choice, here we can also speak of optimism-pessimism levels. With the exception of Y_1

and Z_1, we may term the newly-introduced functions pessimistic or optimistic quasi-conjunctions and quasi-disjunctions of corresponding pessimism-optimism levels.

To illustrate the effect produced on automata collective behaviour by these new levels, we will now turn to considering a model which is a well-known generalization of the allocation game discussed at the beginning of this chapter for homogeneous automata.

Every morning a herdsman has to solve a rather complex optimization problem before sending cattle to pasture, i.e. which pasture should be chosen. He knows of n pastures which are suitable. Other shepherds grazing pasture cattle in the same area are, however, also knowledgeable. It is quite probable that when our herdsman comes to a nice nearby valley he will find that somebody has already got there and his hard work will have gone for nothing. In the highlands he knows the grass may be even scarcer because the summer is dry. In fact, there is another good pasture but it will certainly have to be shared with neighbours and his cattle will eat less than it is necessary.

How can the herdsman achieve his goal, i.e. increase the weight of his cattle? Word has spread that in the adjacent district, the herdsmen were controlling grazing. In his district, however, it exists only in plans. What are local bosses up to? How

are they going to increase the total weight of the livestock in the district?

While the herdsman is brooding over his troubles, we will try to formulate the search for good pasture. The herdsman and his cattle will be modelled by an automaton with n possible actions, in other words, one pasture is chosen from n pastures possible. The automaton has two *a priori* estimates of each pasture: an estimate of the probability with which the cattle will find enough food x_1^i, where i is the pasture number and the average forecast of how many automata may come to a pasture i (x_2^i) at the same time as the automaton under study. These two estimates may be a product of previous experience, knowledge of the quality of the pastures and weather or just empiricism. Roughly speaking, we assume that all the estimates are ternary. Thus $x_1^i = 1$ means that in the ith pasture there is enough grass to feed the cattle, $x_1^i = 0$ means that the ith pasture is poor and, finally, $x_1^i = 0.5$ means that the automaton does not have enough information about the quality of the ith pasture. Similarly, if $x_2^i = 1$, the ith pasture will provide our automaton with enough food even though many automata may be grazing at the same time provided that the resources are divided equally between the automata, $x_2^i = 0$ means that our automaton's share will be too small,

and $x_2^i = 0.5$ shows that more information is required.

It follows that our hungry automaton in choosing the best pasture may benefit on the experience accumulated by the he- sitating bachelor. Computer-aided model- ling has shown that an automaton collec- tive only reaches an optimum (from the viewpoint of local authorities) with certain distributions of pessimism-optimism lev- els. It is noteworthy that when the autom- ata starved, i.e. they failed to acquire sufficient food for several distribution rounds, the proportion of temperate pes- simists tended to grow in time for they appeared to possess a higher survivability than optimists of all levels. The propor- tion of optimists and pessimists in the chain is largely dependent upon the true param- eters of the environment. The same may be said about their distribution with res- pect to levels. In any case, extreme opti- mists and pessimists are not useful in a group, in which they are the first to be- come extinct. On the average, the most stable associations have about 40% objec- tive automata, 40% moderate pessimists, and 20% moderate optimists.

The reason is that in homogeneous groups with neither common fund procedure nor random pair interaction all the automata pass from one state to another in groups. In contrast to this, if a group contains different automata, the pessimists and op-

timists make choices which could never be made by an objective automaton. This helps avoid the spread of states along the automaton chain in groups. We have shown, that the same effect is achieved by introducing a common fund procedure in a homogeneous group of automata playing an allocation game.

In the herdsman model, the players do not deal with the values of environmental parameters themselves but with estimates of them. In one experiment it was assumed that $x_1^i = 1$ provided that the probability of there being enough food in the ith pasture exceeded 0.75. If it was less than 0.25, then $x_1^i = 0$. For all other cases $x_1^i = 0.5$. For the second parameter $x_2^i = 1$ if the ith pasture was attended by less 1/4 of all the herds in the district. When 3/4 or more of the herds are in the ith pasture, $x_2^i = 0$. Otherwise the second parameter was 0.5. These boundaries are obviously artificial, and studies of decision-making in conflict situation show that people resort to many such subjective estimates.

Figure 3.18 gives curves illustrating the attitude of a human player to gains he receives in a game. The values of gains and losses and the subjective estimates of these values by the players are plotted along the abscissa and ordinate respectively. The names given to the players speak for themselves. The analysis of these estimate functions accomplished by Kemeny and Thomp-

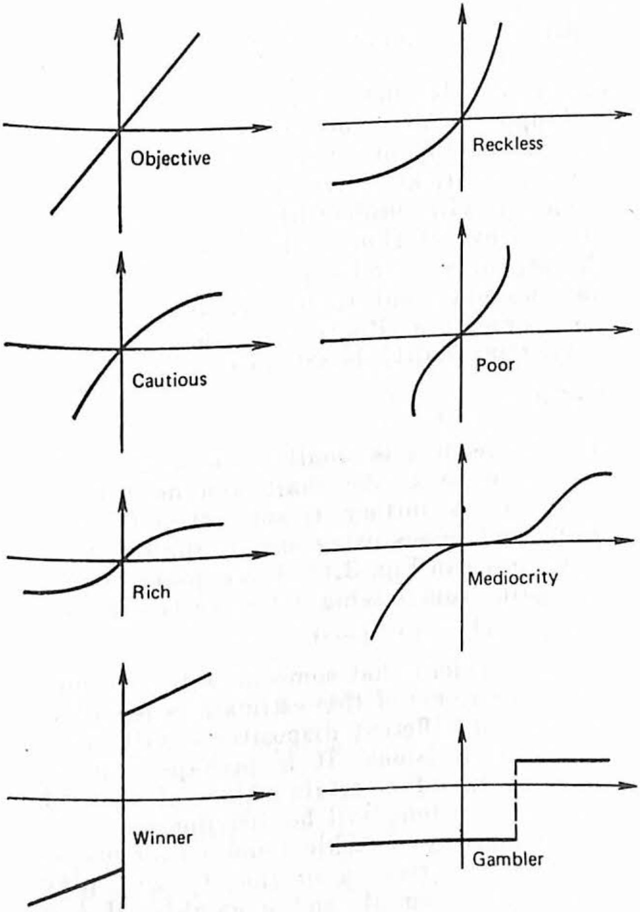
Objective
Reckless
Cautious
Poor
Rich
Mediocrity
Winner
Gambler

FIG. 3.18

son shows that in a group of players with different psychologies their decisions in identical situations vary greatly. Here is one of the models offered by these researchers.

Suppose a Mr X organizes a lottery. He determines the price of the lottery ticket s so that a ticket buyer will win a sum of money l with probability g. The mathematical expectation of the loss suffered by the organizer is $g(-l) + (1 - g)s$. Since he does not want to lose money, he will try to create a situation in which the following inequality is satisfied

$$0 < g < \frac{s}{l+s}.$$

The value of g is small because l is large compared to s. We shall assume that a buyer of a lottery ticket estimates his gains and losses using one of the estimate functions f in Fig. 3.18. The expected value of a gain from buying a lottery ticket is $gf(l) + (1 - g)f(-s)$.

It is evident that someone will only buy a lottery ticket if this estimate is positive. Players of different dispositions will make different decisions. It is perhaps easy to imagine that for certain values of s, l and g their decisions will be distributed in the following way: a reckless and a poor player will take a risk; a mediocrity will play only if l is small, and a gambler if l is larger than the horizontal coordinate of the breakpoint. An objective, cautious, win-

ning and rich player will refuse to play at all a mediocrity will refuse to play when l is large and a gambler will refuse to play if l is smaller than the horizontal coordinate of the breakpoint on his estimate function.

The last two sections suggest the following conclusions. In collective behaviour models, the introduction of inhomogeneity helps to achieve goals which before required additional devices to make the environment induce the members of a homogeneous group to act accordingly. This allows us to assert that the inhomogeneity which occurs so frequently both in nature and in mechanical systems is far from being occasional disruptions of harmony but is instead a reflection of the fundamental fact that the performance quality displayed by inhomogeneous groups in decentralized environments are better than those of homogeneous groups.

3.7. Three More Simple Models

In nature, inhomogeneity is a way of controlling the proportions of different species in biocenoses or phytocenoses. Now we are going to consider two simple models which illustrate this conclusion and which are well-known to students of ecology.

Figure 3.19*a* shows an environment inhabited by bacteria which are drawn as ovals. Some of the bacteria have been

invaded by particles called plasmids. These organic formations are on the border between life and nonlife although self-reproducing, plasmids rely for their metabolism on the

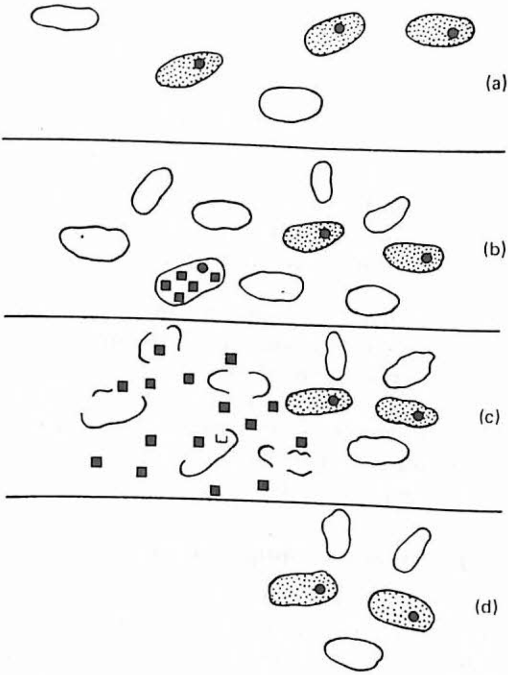

FIG. 3.19

environment, which is the cells of the bacteria. When the environment is not overpopulated and the bacteria have enough food, the plasmids produce a substance

called an immunoprotein. In Fig. 3.19*a* the plasmids are shown as black circles and the immunoprotein as dots. When the number of bacteria has grown beyond a certain population level they starve. The plasmids starve too. As a result the plasmids which are starving begin producing a substance which kills both the bacteria and the plasmids instead of the harmless immunoprotein. In Fig. 3.19*b* you will see one of the plasmids producing this poison which is shown in black squares inside bacteria. By penetrating into the environment, however, the poison destroys all the bacteria in the vicinity that have no immunoprotein (see Fig. 3.19*c*) thus reducing the number of bacteria (see Fig. 3.19*d*). If there are still too many bacteria, then another plasmid falls below the "threshold of life", begins to produce poison and hence further reduces the population of bacteria. Note that the plasmid-bacterium combination only reproduces by fission when there is "too much" food, i.e. more than is necessary for the fission of "pure" bacteria.

This model of a self-stabilizing biotic community may be simulated in the form of an inhomogeneous automaton group existing in environment in which there is a constant supply of food. This food is divided equally among the members of the group. Plasmid-carrying automata are fissionable if the amount of food consumed

exceeds a threshold, here labelled Q_1. The other automata can reproduce by fission at a much lower threshold, denoted Q_2. When a plasmid-carrying automaton receives less food than $Q_3 < Q_2$, it starves and dies destroying all the ordinary automata within a certain range (for example, a torus of cells within a distance five times the size of a cell). To avoid the simultaneous death of all the plasmid-carrying automata, one of them happens to survive. If after this the food level is still less than Q_3, another randomly chosen automaton dies to reduce the population. Computer-assisted modelling of this process revealed a great resemblance between this self-stabilizing process and the one occurring in nature.

The second model for controlling the size of a population is slightly more complex. It is supposed that in a conflict situation (for example, a random pair interaction) the members of a group may resort to two strategies: aggressive or threatening. If the two members in conflict use an aggressive strategy, it is very much like a fight between two cocks or two deerbucks. Both opponents build up their efforts and neither wants to surrender. The winner is thus revealed by the death or running away of one of them. An encounter between a dog and a cat is another interesting example. Here one of them uses an aggressive strategy and the other one a

threatening strategy. When the level of aggression reaches a critical point, the one who was merely threatening seeks to escape. The dog invariably keeps to the aggressive strategy first, while the cat threatens by arching its back and hissing. If the dog is frightened and chooses to adopt a threatening strategy too, then after mutual threat postures each goes its own way. If the dog continues with its aggressive strategy, then the only way for the cat to survive is to retreat.

Both opponents may adopt threatening strategies from the very beginning. They take different ritual threat postures until one of them surrenders and shows this by taking up a special submissive posture. Such rivalry exists among dogs, grey geese, and many other animals.

Consider a model of such rivalry. Let A and T denote the aggressive and threatening strategies respectively. Table 3.6 shows the estimates of all possible combinations of encounter.

Table 3.6

		Second	
		A	T
First	A	(−5, −5)	(+10, 0)
	T	(0, +10)	(+2, +2)

The pairs of figures in the table are arbitrary estimates of the gains and losses suffered depending on the strategy chosen. If one opponent (say, the first) chooses strategy A and the second strategy T, then the former gets a return of 10 units while the latter gets nothing. What cause might have this effect? Initially we might say that a victory in a conflict is a gain of $+10$, a serious injury or death resulting from a buildup of efforts in the strategies A is equal to (-20). Since the outcome of a fight between two aggressive opponents is equiprobable, the expected value of a reward or a fine with the pair of strategies (A, A) is $0.5 \times 10 + 0.5 \times (-20) = -5$. Similarly, in a fight with strategies (T, T) the expected return is $0.5 \times 10 + (-3) = 2$. Here the estimate (-3) is due to the nervous strain inevitable in a long conflict with strategy T. This strategy exhausts the animal's nervous and other resorces. Hence, Table 3.6 is a game matrix.

Consider an organism which can change its strategy depending on the circumstances entirely on its own. It can be modelled by an automaton with two states corresponding to the strategies A and T which are chosen with respective probabilities P_A and P_T, such that $P_A + P_T = 1$. Consider a collective of such automata, assuming it is inhomogeneous. The inhomogeneity is described by different values of P_A. For instance if $P_A = 1$, the automa-

ton is a pure aggressor. In all cases it keeps to the strategy A. If $P_A = 0$, it always chooses the strategy T.

As in the previous model, we specify certain thresholds Q_1 and Q_2. If an automaton gets a gain exceeding Q_1, it starts "breeding", yielding two automata with the same value of P_A each. If its absolute fine exceeds Q_2, the automaton dies. What then is the optimal value of P_A in random pair interaction between the automata in the group? Computer-assisted modelling has shown that if there is a large number of automata and the P_A distribution is close to uniform, the group evolves into a homogeneous collective of individuals whose P_A approaches 8/13. It follows from game theory that a mixed strategy in which the strategies A and T are chosen with probabilities 8/13 and 5/13 is in a sense the best for a player. It maximizes the possible assured gain when his opponent's actions are least profitable to him. It would be of interest to obtain experimental data resulting from the observation of animals (for example, cats) and showing how often they choose the strategies A and T when they have to deal with an equally strong opponent. It is a pity that we don't have such data.

Now we return for a moment to the risky business of Lieutenant Schmidt's children. The Sukharev Tower Pact inspired us to study a number of models of collective in-

teraction and agreement. Had the legendary hero's children been more educated in the field of decentralized control, they could have made much more money than they really did. This stresses the practical value of these models.

In conclusion we would like to offer you one more model which seems applicable to the method chosen by Lieutenant's unscrupulous "children" to earn their daily bread. Suppose their operations expand so as to include areas where none of them has ever been. Such areas would appear to be equally valuable and their distribution would not arouse debate. This makes drawing lots a formality. Now the children go to their districts and get down to business. After a while each of them can estimate the average profit which may be obtained in a particular district. Justice demands that at the next conference the profits are equalized with the help of smart-money or a common fund. Greedy as they are, the owners of the richer districts will not want to contribute. So they will lie. If he is talking about the average profit in his district with a colleague whose profit is higher or who is unlikely to forward any claims to the fund, a cunning "son" of Lieutenant Schmidt can afford to be frank. However, he will lie to all the rest reducing the real figure to a level which makes his district a matter of no interest to all those present. In this model, the members of the group

must be informed of the real profits of the other participants.

Here an increase in the volume of information improves the situation in which an automaton has to function. We would like to draw your attention to this significant trait of the model discussed. There is no direct dependence in a reflexive behaviour model.

All the models offered in this chapter have one important feature. Since we have considered an automata collective as a model of a biological, social, or technological system, it is easy to see that the system functions in parallel mode, i.e. all its subsystems operate independently of one another. They don't have to wait for the effects produced by other subsystems. In practice, this is rather a rare case. In complex systems, the operation of subsystems is interrelated and takes place at times according to the relations within the system. These relations may be either probabilistic or deterministic. For this reason in the next two chapters we intend to study decentralized control when these additional restrictions are enforced.

Chapter 4

Jump the Queue and Call It Fair!

Dura lex, sed lex.

4.1. Where Do All the Queues Come From?

Ostap Bender, the great schemer and hero of *Diamonds To Sit On*, another famous novel by Ilf and Petrov, didn't like queues. He considered queueing up for something either a luxury he couldn't afford or a personal insult. Pushing his way through a crowd to the reception-room of the manager, executive, or social worker he was going to make a fool of, Ostap would cry: "I'll only take a minute! Don't you see I haven't even taken off my galoshes!" This reasoning was little comfort for those waiting patiently in line. Obviously, he considered the last argument convincing enough to give him the right to jump the queue.

Unlike Ostap Bender, however, we are more interested in his first argument, i.e. that he was not going to occupy the manager's attention for a long time. Would Ostap have been right in thinking that it was fair to jump a queue if he wasn't going to take up too much time? How should a waiting line behave in a situation like this? And, finally, what could the manager do to handle the situation?

The queue! In life it has become as essential and inevitable as meals, sleep, and

entertainment. We queue for a meal in a good restaurant or for a show at a box-office. At times, we are inclined to believe that a queue is the result of somebody's wickedness or wrongdoing. In fact, the queue is as natural as a snowfall in winter or a rainfall in summer. Poor management does not create a waiting line, but it can make it longer.

A waiting line obviously implies the presence of those who are waiting. Note that a waiting line need not consist of people, it may be, cows waiting to be milked. Things also queue up. For example, radios wait for being repaired, deposits of mineral resources wait for being explored, while those that have been explored wait for being utilized. A waiting line may also be composed of immaterial objects: a new scientific idea, for example, waits for being developed and applied.

We are going to term any object in a queue, regardless of its nature, a user. Note that a queue does not mean any association of users, it only means a group linked together by a common goal, the goal being the desire to get service.

Service need not be necessarily active. Suppose at 7.58 in the morning you are in a queue waiting to sign on at work. You put your signature yourself, but from the viewpoint of the queue you are served by the register.

We define a *service channel* as the set of

things required for the service, such as service personnel, a service site, rules of service, etc.*

A queue embraces the whole set of users, channels, and rules of interaction between the users and the channels.

To study a queue we must know how users enter the queueing system, what mechanisms or models of mechanisms cause the users to desire to go into service, what are the characteristics of the service procedure, how the queueing users' behaviour and interaction are or may be organized, how the users get to a service channel from a queue, and how the service channels get their users. Finally, we must have some information on all the possible mechanisms of interaction between the different service channels when the queue is obtaining and serving the users.

A meticulous reader may catch us on a point of the validity of our definitions. On the one hand, we regard users as an integral element of the queue and, on the other hand, we say that they come into the queueing system from outside. This contradiction is clearly just a seeming one. Indeed, users come from outside to become elements of the queue. The same thing

* For someone familiar with the terminology of queue theory a service channel is the equivalent of what in the classical models of the theory is known as a server. Our treatment of a queueing model, however, is broader.

happens to firewood arriving at a heating system.

In order to organize control in a queue, we must be able to assess the rules of the queueing system behaviour as well as to define its performance quality criteria. It is perhaps easy to see that these criteria are of contradictory nature. As a rule, a highly cost-effective service system is not the same as one with a high performance quality. City dwellers would hardly be enthusiastic if local authorities were to decide to make the municipal transport system more cost-effective by increasing the number of passengers carried in one lift. It is the contradiction between different performance quality criteria that makes a queue a fascinating object for study from the viewpoint of the organization of optimal control.

How do users come into the queue? The simplest way for them to arrive is in uniformly spaced intervals of time. If the service time is shorter than or equal to the intervening interval, there will be no waiting line at all. If the service time is longer than the interval between customers, the line will increase infinitely.

A channel's *capacity* is the maximum number of users which may be served by a channel per unit of time. The channel *load* is a fraction of time during which a queue is busy. If the time between user arrivals or, as we shall sometimes put it,

service requests per unit of time is constant and the service time is constant, there will be no waiting line provided the channel load does not exceed the channel capacity.

If the rate of arrivals or time in service are subject to random variations, a queue is inevitable. This is true even if the channel capacity exceeds its load. The greater the variance in the interarrival interval and the period in service, the longer is the waiting line. The queue also increases as the channel load approaches the channel capacity. The result is an infinite increase in the waiting line. Having established the relations between the parameters of a queue and the queueing system, we are now better equipped for trying to change them and thus reduce the waiting line.

Once, when working on Chapter 3, we decided to have a break and go to the cinema. We tried to call the nearest one to find out what was on. The line, however, was constantly busy. After twenty minutes we at last connected and were treated to a recording which gave us all information we needed. The voice on the phone, however, greeted us nicely, reminded us which particular cinema we had phoned and invited us to visit it. We looked at our watches. It was 4 p.m. but the voice had cheerfully informed us of the films we had already missed. We decided that if it had taken us less than twenty minutes to get in touch with the cinema, we would have

excused the voice for not being so polite. Then we took our pencils.

We discovered that a reduction in the length of the spoken text would increase the channel capacity thus changing the capacity/load ratio to the extent that the average waiting time would be reduced nearly five-fold. It would have taken us only four minutes to get the cinema on the phone. A 16-minute gain would have been a sufficient reward for having to change the tape in the answering machine after every film. Suppose we were just unlucky and that it might have taken someone else ten instead of twenty minutes to connect to the answering machine. Even so the reduction in service time by two would give an hourly economy of about ten man-hours of dialing. The situation cannot be remedied because the thousands of man-hours a month lost in futile attempts to call to cinemas and the resulting additional load on telephone channels and commutations equipment have no impact whatever on the performance quality criteria of these cinemas and their information service. The same may be said about the great majority of telephone serving systems which are barely accessible to the user.

It is obvious that a queue causes moral and material damage. Is it possible that this damage is less important when we speak of inanimate objects waiting in a queue? Can we assess the damage?

First, any queue, whatever the user, occupies space. The longer the waiting line is, the larger the storehouse needed to accommodate it. It should also be taken into account that, roughly speaking, more than a half of waiting lines will be longer than average and will take up much more room. The stores housing waiting lines (for example, warehouses for queueing articles or buffer storages for data files waiting for being processed) make queues expensive. In a number of cases it is cost-effective to reduce the space occupied by the waiting lines by introducing additional service channels.

Second, users in a line are taken out of use. People queueing up in a shop are not working, reading, or playing with their children. Trucks waiting for being repaired are not carrying cargo. The cost of what makes up a queue is part of the cost of the system. If a large oil trunk line is full of oil, this oil is out of use. It becomes a part of the pipeline and its cost is included in the cost of the line. The same thing happens in a queue. Parts waiting to be machined are out of use and in the average length of the waiting line must be included in the cost of the machining system. The average number of automobiles waiting for being repaired is an integral part of an automobile service system and thus their own cost is added to the cost of the system.

It follows that a reduction in the average length of a queue is obviously economic.

The average length of a queue may be reduced by reducing the service time or, which is actually the same, by increasing the serving systems capacity. Frequently, however, a change in this parameter is beyond our power. Is it still possible to cut the average length of a queue?

We have already pointed out that, among other things, a queue is characterized by the relations between those who make up the line, in other words, an agreement between the users concerning their behaviour in the queue. We will call this agreement a *service discipline*, and there are several.

The most commonly-known discipline is called "first-come-first-served". This is an ordinary queue. There are also exotic disciplines like "last-come-first-served", which makes sense for different reasons including the need for building a store housing the line. Such a store is called a magazine similarly to a magazine feeding cartridges to an automatic rifle. It functions like a stack in a computer. "Last-come-first-served" discipline is often used in air defense systems. Here the last user (a hostile aircraft) to enter the service area has more chance of being served, i.e. shot down, because it stays in the service area longer than the other users, i.e. other aircraft.

There also exist priority service discip-

lines which allow a priority user to jump the queue. A typical example of an assigned priority is when women and children are pushed first in emergencies. Priorities are many and diverse but they all depend on the state of the queue. If we can't change the system load or the service channel parameters, i.e. their capacity, a priority discipline is the only way of interfering in the operation of a queue, i.e. controlling the queue's behaviour. Is it possible to improve queue performance quality by introducing priorities? If it is, then which parameters may be improved and what priorities should be introduced? Are there any methods for organizing the collective behaviour of the users or service channels to help create a system of priorities which optimizes queue performance?

4.2. Barbers, Clients, and Priorities

Suppose a waiting line in a shop consists of five men who are going to make similar purchases. It takes the salesman approximately the same time, say, about 6 minutes to serve each of them. A young man enters the shop to buy a box of cigarettes and takes his place in the queue. His service time is 30 seconds. Now let's see how the action unfolds beginning from this moment. The salesman starts to serve the first man in the queue whose waiting time

is equal to zero. The second man has to wait for 6 minutes, the third 12 minutes, the fourth 18 minutes, the fifth 24 minutes and, finally, the young man who has dropped in for a box of cigarettes has to wait for half an hour. The total waiting time of all the buyers is an hour and a half. That is 15 minutes per buyer on the average. Keeping this figure in mind, we proceed. Suppose the impatient young smoker buys his cigarettes out of turn. Instead of having to wait for 30 minutes he doesn't have to wait at all. The corresponding waiting times of the other five men in the queue are 0.5, 6.5, 12.5, 18.5, 24.5 minutes. Now the total waiting time is 62.5 minutes and mean waiting time is 10.4 minutes, i.e. it is nearly one and a half times as short. The average length of the queue during the period discussed has reduced from 2.4 to 2.

It is obvious that our example does not give a completely realistic picture because the young man was the last client to join the queue. Yet it shows one mechanism which reduces the queue size and the mean waiting time by granting highest priority to the user with the shortest service time. So Ostap Bender appears to have been right in demanding that he be served out of turn because he would "only be a minute".

If you have ever had to waste your precious time in a queue, you know very well that nothing is so infuriating as when

somebody tries to jump the queue thus increasing your own waiting time, even if you know that this act of self-sacrifice on your part will improve the queueing system's performance.

If it is difficult to imagine how a system of priorities could be introduced in a queue, it is perhaps wise to think whether priorities may be imposed on a queue by service channels. However, it is necessary that queue discipline implies the right of the channels to assign priorities or, which is actually the same, to choose users from the waiting line themselves.

Now we leave the shop and go to the barber's. If there are several barbers we have a queue with several equal capacity channels. The queue consists of clients belonging to different categories which differ in service times: some want a shave, some want their hair cut in a simple or stylish fashion, others want a hair cut, a shave and a massage, etc. We have already pointed out, that in order to reduce the mean queue size it is expedient to let clients with the shortest service time go to the head of the queue. Were the users' service times known *a priori*, the management could compile a list of priorities based on the principle "fastest-served-first-served". In a situation like this, Ostap's cry "I'll only be a minute!" gives him the right to jump the queue. (The additional argument concerning the galoshes remains far beyond the scope of

our model.) The list of priorities may be simple enough. For instance, "those wanting shaves served first" will do.

Things are more complicated when *a priori* information concerning the service times is not available, when these values are random or change over time. In this case our aim is to assign dynamic priorities to the queue, i.e. as the system operates. In our approach the priority system must arise from the collective behaviour of the service channels.

The solution of a collective behaviour problem demands that we formulate individual preferences in a manner which ensures that the individual needs of an association of users are satisfied as far as possible and the required performance enhancement at the queueing system level is achieved. We have already pointed out that it is hardly possible to talk the user into believing that his waiting time has been increased for his own sake. It is relatively easier to persuade the service channel.

Suppose we assert (reasonably) that a barber's goal is to maximize his fee. Then if he thinks that a client with a shorter service time is more profitable, he will try to handle the situation accordingly and take him out of turn. You often see a barber invite an old client out of turn. Other men waiting in line may protest but we have already agreed that a barber has the right to have favourites. Before inviting the

first man in the queue, for example, the barber may announce that he serves students out of turn. As was pointed out above, a less time-consuming client must be profitable to the barber. This may be achieved by the introduction of a fee to be paid for service irrespective of its length. Generally speaking, a constant fee leads to longer service times. Money paid, do what you are paid for. With this in mind, we may introduce an initial fee which is independent of service time, like the one in a taxi.

We began this discussion by saying that queues cover a far greater variety of systems than ones in which the users or the channels are human beings. For example, various technologies such as communication systems, computers, and transportation systems, (e.g. the conveyer belts feeding coal to power station's input hoppers) are all queues. For this reason, a study of the behaviour of channels and users in such systems demands that we formalize their behaviour so as to simulate it using simpler devices. This approach will then permit us to build and examine the models of collective behaviour in service systems. If in reality local goals are achieved with reliance on more ingenious methods, so much the better for the queueing system. For convenience, however, we will have to be content with the terminology of a barber's shop.

Now we classify users by type. Every type

of user needs a service time lying within a certain time interval. Every type is assigned a number common for all the users of this type and termed a user number. Regardless of the number, all users pay the same fee prior to getting into a service channel. Suppose this fee consists of K identical coins placed by the barber into a special money-box. If he services the user for T units of time, the barber takes T coins out of the money-box, i.e. the amount of money in the box is reduced by one per unit time. If $K \leqslant T$, no more money is left in the box and the user leaves for good. If $K > T$, there are $(K - T)$ units of money left in the box, so the barber notes the user's number on the money-box, and this gives the user priority next time he comes.

If the mean waiting time in the system is K, priority is given to users whose service times are shorter than the mean service time in the system. If a priority-holder comes for servicing again, his entrance fee will be added to the money still in the box bearing his number. Thus, over a sufficiently long period of time all users whose service times are less than the mean in the system will be arranged according to the amount of money left in the money-box, i.e. according to the mean value of their service times in relation to their service frequencies. For instance, if A-type clients give the barber one dollar of profit and come to him ten times a day, they are bet-

ter than B-type clients who give one hundred dollars but come to the barber once a month. If all the types of user are arranged and their priorities numbered, then in order to optimize the queue size B-type users must be given a lower priority than A-type users, however small the effect this arrangement brings.

This priority assignment technique differs but slightly from the collection of current statistical information concerning users' characteristics and the resulting optimal priority system. We can't be content with this solution. First, the technical complexity of the queue increases as the number of user types. The more users go into service, the more money-boxes must be available. Second, the way the problem is stated eliminates the very idea of decentralized control because it implies that the priority system is optimized in each particular service channel. We will try to simplify the system and thus perhaps reveal some new fascinating traits in its behaviour.

We assert that the number of priorities in each channel is limited. Suppose each barber has only a few favourites. In the simplest case this number is two. Now the barber needs two money-boxes only. The first two profitable clients will have priority in this channel and at this point the establishment of a priority system will come to the end. To avoid deadlock, let us introduce a kind of competition between the

users in the channel. If a non-priority user arrives at a channel, the contents of the money-money-boxes are compared and the money-box with the smallest sum is emptied so as to begin to work on the newcomer. It is natural that we want to know how many money-boxes are required. The smallest number is two, and the largest should not exceed the number of user types. A study of behaviour of such a system has shown that this increase in the number of money-boxes leads to insignificant changes in performance quality which are incomparable to the cost caused by complicating the control system.

The introduction of such a system demands that one more question be answered. If a money-box is captured by a user with comparatively small (but obviously not the minimum) mean service time and if this user manages to save a handsome sum of money in his money-box, it is unlikely that other users with money-boxes will be able to leave him behind, for his savings will grow at a slow but steady rate. If, at some moment, this user quits for good, it is highly probable that the channel will still retain the priority involving the now non-existing user. This will cause the whole system to malfunction. There is no doubt that we can improve the situation by raising the rules of priority assignment to a higher "intellectual level", but this is not what we are after. Another try should be made to

find a simple set of rules of interaction and to obtain optimization algorithms. It does not follow from here that a wise barber must bury his talent. The best way to handle the situation is to restrict the volume of the money-box. When it is full, the extra money is given to the cashier. The introduction of this restriction appears to improve a system's behaviour when the users' characteristics vary with time. There is an optimal money-box capacity minimizing mean queue size for each particular degree of the system's non-stationarity, i.e. for each average time interval between the changes in the characteristics of the users. The higher is the rate of change, the smaller is the money-box capacity.

We hope this reminds the reader of the fact that expedient automata operating in alternating random environments have an optimal memory capacity. The analogy is complete, because the smaller the money-box capacity, the more teachable and pliable the system becomes. On the other hand, a more flexible system is more liable to make an occasional mistake by failing to see the difference between profitable users. The optimal capacity is actually a compromise between the quality of decision-making and the time necessary to decide.

The assignment of different priorities for a variety of channels improves another of the system's parameters, i.e. the number of transitions. We define a transition as a

change from one type of user to another type of user. If two users of a single type are served one after another, there is no transition. It is perhaps easy to understand that if a particular channel gives priority to users of a certain type, the majority of users belonging to this type are sure to require service from this channel, thus a more uniform flow of users results in a lower transition rate.

What good is this? First, a transition always means a loss. It may be readjustment of equipment, additional disturbances caused by a commutation in a communication channel, or a time loss.

Second, a lower transition rate tailors the service channel to fit users of a particular type and thus reduces the service time of the priority users.

Now we have touched upon a point which promotes a broader treatment of our model. One and the same user may obviously have different service times in different channels, while the service time in the same channel may vary depending on the particular user in it. In this case a redistribution of users among the channels changes both mean queue size and the mean queue capacity.

4.3. How to Learn to Be a Foreman

Consider a workshop consisting of several interchangeable work areas. In each area a worker carries out several operations

which make up an order. Irrespective of whether the work is done by a single worker or a team of workers, we will use the word worker to refer to the combination of the worker(s) (the service channel) and the work area.

To fulfil a work order, i.e. to serve a user, it is necessary to complete a series of operations. Suppose there is a work order for a device in a tool shop. The set of operations may include milling, drilling, polishing, scraping, assembling, etc., but the work order will be characterized by the labour input and the complexity of operations. The labour input to an operation is determined by the time assigned to do it and its complexity is determined by the required qualification of the worker. The amount of money to be paid for making the device is determined by the amount and complexity of labour.

Each worker has different qualification levels for different kinds of work. His qualification to do a particular operation is characterized by labour coefficient for each type of work.

It is easy to see that it will take different workers different periods of time to do all the operations specified in a work order. So the time spent will depend both on the set of time assigned to the operations and the corresponding set of labour coefficients for these operations. It is also obvious that the labour productivity in the workshop

is dependent on the distribution of work orders among the workers. The more closely the working hours needed for the separate operations are made to correspond to each worker's individual qualifications, the higher the average labour productivity will be.

The main problem a foreman faces is to organize the production process, i.e. to distribute the work orders among the workers using his knowledge of the individual abilities of each worker and the characteristics of the work orders which make up the production goal of his workshop. An experienced foreman who knows both his workers and their abilities and the production standards usually copes with this problem.

Once in a while, however, even an experienced foreman can make a mistake. For some subjective reason he may either underestimate or overestimate the abilities of a worker. The fact that some of his subordinates have learnt how to do a certain job better and some have forgotten how to do it may escape his attention. The foreman's right for personal decision-making may lead to conflicts breaking the workers into favoured ones and offended ones. It may often be the case when certain jobs became profitable or unprofitable due to unsatisfactory rate fixing, unequal intensity rates and the inadequate system of payment, etc. When we studied this model in a real plant, we came across the fact

that the same worker can make either 11.5 or 1.5 roubles a day depending on the assignment he was given by the foreman.

Note again that the situation described is typical of a great number of purely technical systems; we chose to give you an example with people merely to help writing and reading about it.

In order to optimize the operation of the workshop, we may keep the foreman away from distributing the work orders and instead allow the workers in the group to do this. It is clear, however, that no decision-making procedure based on auction or voting will be effective in a production system and for a number of reasons.

Generally speaking, the work distribution could be optimized using known mathematical methods, provided all the parameters of the workshop's work plan and the workers' qualifications are known in advance and are not liable to change during the period being studied. The trouble is that work orders do not arrive in a uniform flow due to factors outside our control. A worker's efficiency is also subject to occasional changes caused by his mood, health, and fatigue. In the time interval being studied the production goals may also vary. Hence, any rigid distribution of work will yield situations in which some of the workers are idle, while others will have to deal with a queue of urgent orders. The solution, however, is within reach. The distribution

of work orders among the workers may be operationally controlled by introducing a dynamic priority system into the queue of orders.

The above model suggests the following work-rules for the interaction between the workers which ensures the decentralized distribution of work orders among them.

All the work orders are divided into groups according to the labour needed at each stage in the work. At regular intervals, say, every morning or every Monday morning, each worker may announce the one or two types of work he would like to do. Indeed, he may do it upon completion of his previous work instead of at regular intervals. His request is satisfied provided such an order is available. All the work orders make up a queue, which is organized according to their importance and the deadlines set for the completion of the work. If the work announced by the worker is out of the queue, he gets the first order in the waiting line.

Technical means required for the realization of this idea include a box where the orders are stored as they arrive and a clerk who classifies them according to the above criteria and assigns them to workers using the algorithm we have just discussed. To improve the performance quality of the system additional information may be written on a special notice board, where workers would assign numbered

priorities to work orders so that everybody could see them.

It is logical to assert that the choice of priorities would depend on each worker's personal likes and dislikes. We may also assume that this choice would reflect his idea of what is or is not profitable to him. We think it is evident that in this case the growth of the total wage would correspond to a rise in efficiency.

In such a system, there is a danger that all the workers will unanimously choose one and the same work order as being "most profitable" and so the priority system will not come up to our expectations. It is noteworthy, however, that the more workers keep to the same type of work, the less profitable that sort of work would be for the average worker, who will have a chance to get it comparatively seldom. It is readily seen that two roubles a day extra are more attractive than twenty roubles a month extra. There is no doubt that those who take part in the distribution of the orders will be good enough at arithmetic. We may hope that priority orders will include those which bring a steady profit on a regular basis rather than those which lead to a dramatic but occasional rise in the wages.

Giving assignments to workers according to priorities in a fashion similar to the model described in the previous section will stimulate workers to acquire specialized

skills which, in turn, will improve labour efficiency and the capacity of the workshop and will make the priority types of work ever more popular.

It is not hard to see that the organization of the operations' distribution similar to the one described above can meet with difficulties including those of a psychological nature. No experiment in a real production system has so far been conducted in this country though the necessary preparations were once made in several enterprises including the *Kalinin Electrical Engineering Plant* in Tallin. During preparations, a distribution system was studied with the help of a model.

Now let us look at the construction of such a model. Initial data were obtained in two workshops at the *Krasni Proletari* plant in Moscow and one workshop of the *Pneumatika* plant in Leningrad involving 22 parts and 3 workers, 47 parts and 5 workers, and 25 parts and 12 workers respectively. Computer was used to model a process in which the parts arrived at the shops to be finished later. Each time a new work order had to be assigned, the computer requested assistance. The operator had experience in distributing work and knew all there was to know about the work orders and the workers. In the model, this man was an objective, well-informed foreman. In this manner, the shop's operation was simulated and studied for 500 hours. Then the distri-

bution procedure was realized as a result of collective behaviour. Each worker was simulated by two money-boxes of the type described in Sec. 4.2. The simulation of 10,000 hours of workshop's operation revealed that for real data the introduction of priorities increased the capacity of the shop in comparison with "first-come-first-served" discipline by 3-7% depending on the shop. Note that a well-informed and objective foreman in a situation like this achieves practically the same results as a collective of workers simulated by a rather primitive local decision-making procedure.

4.4. One Circus Ring Is Not Enough

The circus is perhaps the oldest or, if not, one of the oldest crafts man has ever known. Surely, you might think today's circus differs greatly from what it used to be thousands of years ago. However, since a round circus ring was first introduced its diameter has remained the same in circuses throughout the world. Every act on a circus programme is rehearsed to fit the ring size. A risky acrobatic feat (are there any that are not risky?) would be fatal should you change a ring diameter. The small size causes a contradiction between the diameter of a circus ring and the desire to sell more tickets. If the ring diameter cannot be varied, how can you get more people to watch the show? The unlucky ones

sitting too far from the ring will be unable to enjoy the show. If the big top is bigger than it should be, the distant spectators will have no sensation of being personally involved, a feeling the circus is so famous for.

This problem was solved by having several rings under one tent. The circus show could then become a non-stop succession of acts either running in parallel or, if need be, repeated sequentially in the different rings. Thus the circus "capacity" was increased and the spectator could see a more action-packed show with no time-consuming interval to prepare the ring and the performers.

Soon the circus manager had to face another problem. The variety of acts on the programme had somehow to be arranged on all the rings available. This would not be difficult if the programme was well-timed, i.e. if the time required for the preparation and performance of a particular act is known in advance and with precision. However, in a circus show, especially the one starring guest artists, it is next to impossible to accurately predict the sequence of acts in each ring and bring the show to a desired time limit.

This problem exemplifies a well-known class of problems of optimal queue scheduling with time periods necessary to readjust the equipment. The classical example is the scheduling of machines

which make intricate products but need some time for readjustment between operations. As in previous examples, we went into a discussion about circuses with many rings merely to have a vivid picture of the model we wish to consider below.

The model we are going to study is far more complex and serious than the one used by a director of a circus show. Yet the method we are offering may as well be used to schedule acts in the many rings of a circus. If this book is ever read by a circus director and he should venture into using this method, we are sure he will never regret it.

A network of computer systems working in economy, science, and information service contains a group of systems with a permanent set of programs stored in their memories. These programs are designed for the solution of a finite set of problems $N_1, N_2, \ldots, N_k$. The number of computers in the system is equal to l with $l > k$. Let us label the programs intended for solving some problems by $M_1, M_2, \ldots, M_k$. These programs are run to solve problems as required by the users. The requests for the programs come to the input in a random fashion and the system has no *a priori* information concerning the characteristics of the flow.

Each computer in the system may execute a certain program M_i ($i = 1, 2, \ldots k$). This means that the computer's immediate

memory stores a program for solving the N_i problem and stores all the necessary initial data. Like a circus ring ready for a certain group of performers, the computer is prepared for the execution of a particular program. The computer system's adjustment is decentralized by giving each computer an automaton having k states. Let p_{ij} be the probability of a change from state i to state j (the change from a computer running the M_i program to that running the M_j program). As always, $\sum\limits_{ij} p_{ij} = 1$. If automaton A_m (m being the number of the computer in the system, i.e. $m = 1, 2, \ldots, l$) is set to execute the M_i program and is idle until a request for that program's execution comes to the input, then A_m adopts the request and receives a reward signal. A_m is an automaton with variable structure, as described in Chap. 2. Having received a reward signal, it raises the probability p_{ii} and proportionally reduces the remaining probabilities p_{ij} for $i \neq j$. If the automaton, as the request for solving the N_i problem arrives, could handle the program but at the moment is busy dealing with the same problem as requested earlier, it is also rewarded and changes its transition probabilities in the same way as it would if it were to start serving the newly-supplied request immediately.

Suppose the A_m automaton is ready to

execute the M_i program and is idle while it receives a request to execute the M_j program (with $j \neq i$), then A_m is either fined or rewarded depending on the particular situation in the computer system. If there are idle automata ready and waiting to execute M_j, they accept the request for service and no signal is received by A_m. It keeps waiting for an appropriate request. If there are no automata ready and waiting for M_j in the system, they refuse to serve the request and all the idle automata, including A_m, are fined. This fine makes A_m reduce the value of p_{ii} and proportionally increase all the other values of p_{ij} with $i \neq j$.

What are the benefits of this model of adjustment? It is not hard to see that a computer system involving adjustable automata search the input flow for the service requests which are the most common in this flow. Coming back to the barber's shop for a moment, we can say that the barbers are always ready to serve a frequent visitor and to ask a newcomer to wait as they are expecting clients who are coming in a minute.

The owner of such a barber's shop would hardly like the idea of idle barbers and unserved clients. The same is true of the designers of a computer system. Idle computers generate only loss and are fined, and the fines are to be paid.

A possible remedy is initially to let only one computer adapt itself to the coming

flow of service requests. When it is adapted, the remaining flow is then used to teach the next computer, and so on. Then some nth computer will have to deal with a "tiny brook" of rare requests. This computer may be supplied with a buffer in which the rare requests queue in a line of a tolerable length instead of being offended by refusing to serve them at all. In a barber's shop the role of this computer would be played by a beginner who serves a separate queue of occasional customers. Readers who frequent beauty-salons must have seen this procedure realized more than once.

The difference between these two methods of automata adaptation may be illustrated by an example obtained by computer simulation. We will assess a system's performance quality by the ratio $H = \mathscr{L}^*/\mathscr{L}_t$, where $\mathscr{L}^*$ is the number of served requests during a certain time interval and $\mathscr{L}_t$ is the total number of arriving requests in the period. Suppose the computer system consists of two computers adapted for the execution of one of four programs M_1, M_2, M_3, and M_4 on either of the two automata A_1 and A_2. The system receives a flow of requests for executing these programs. In the experiment, the characteristics of this flow were unknown to A_1 and A_2 and were prescribed by the requests which arrived with the probabilities: $P_1 = 0.15$, $P_2 = 0.30$, $P_3 = 0.45$ and $P_4 = 0.10$. Before the system could adapt $H = 0.5$. After

it had adapted to the flow using the first method discussed, $H = 0.54$. It is easy to see that the method is rather inefficient. If the second automaton deals with the flow remaining after the first automaton, which is tuned to the execution of only one program, has selected requests, then, the second method being applied, $H = 0.57$ after A_1 has been adapted, and $H = 0.63$ after A_2 has been taught too. If the second automaton does not lose requests, then for a certain period of time $H = 0.85$.

This distribution of the computers in a computer system among service requests which arrive in a random and unknown fashion is like many of the technical problems arising in the control of complex systems. The technique described in the last three sections of this book is applicable to the control of channel commutation in a communication centre, pump switching in a large water-supply line, operation of a gravity hump in a railway sorting yard, etc.

4.5. Problem Faced by Housing Board and Similar Problems

Our discussion of decentralized control and collective behaviour has so far implied a kind of behaviour intended to satisfy some utility criteria. Commonly, however, the purpose of behaviour at a system level is to achieve a *coordinated behaviour* of the system's elementary compo-

nents. The system's behaviour is organized only when these components can agree with each other.

In the previous chapter, we dwelt upon such a possibility, and looked at the common fund procedure and the rules for random pair interactions. On the other hand, obtaining agreement may itself be the purpose of a system's behaviour. In a sense, agreement-seeking behaviour is a higher level of control than those we have described so far. Here decentralization means that agreements are not imposed from outside but arise out of the interaction of the objects.

At the top of the hierarchy of levels defining the behavioural rules for the lower level, the rules must be simple and cannot come about by any other mechanism but random exhaustion. However limited and directed, it is still exhaustion.

The problem of agreement is usually complicated by the conflicting interests of the contracting parties. If a man and a woman are in love, the road to the altar will have no insurmountable obstacles. An attempt to come to agreement when dividing the property after a divorce, however, is often futile until the case is brought to court. The centralized solution of the problem by the judge reduces the total value of the property that is divided by the court costs.

Consider several situations in which to

tions in France. When the latter mechanism is at work, the final decision may be opposed by the majority of voters. On the other hand, the right to veto seems to make balloting a totally impracticable decision-making technique.

Yet we are going to make an attempt to present an efficient method of arriving at a decision in a situation when it may be vetoed or confronted by the obviously conflicting interests of those taking part in the vote. For this purpose let us turn to the problem of a housing board as formulated by M. Tsetlin. In a lecture delivered at a session of the Physiological Society on 23 February 1965 in Moscow, he said:

> "In a nutshell, I will explain why it is difficult to divide the available accomodation among those who need it and how it is connected with automata. Frankly, there is not as much accomodation as those wanting it would like there to be. If we had more flats than requests, there would be no housing board or it would be out of work. Let us assume that all the flats to be divided are equivalent. Otherwise, we would face a variety of problems, i.e. how to divide one-room flats, two-room flats, etc.
>
> "Suppose we have N requests and m members of the housing board, m being rather small. Now let us look at the board at work. Every board

member takes a list of requests and scrutinizes it to see who really needs better housing, then the next most in need and so on, until he compiles a list of requests in the descending order of priority. If you excuse my labeling people with letters, the list of requests would be

$$\alpha_1,\ \alpha_2,\ \alpha_3 \mid \alpha_4,\ \ldots,\ \alpha_N,$$

where the vertical line shows the number of flats so far available which can be distributed. It should be emphasized that the board member will select the names of those to the left of this line with special care. The same will be done by all the rest. When they are through, they compare lists or perhaps write columns of names on a blackboard. Then the only thing they can do is to shrug their shoulders and scratch the backs of their heads. Here balloting is useless. We have fewer flats than we have requests and the lists will all be different. Suppose I am one of the voters. I see one of my choices crossed out and refuse to vote at all. I am prepared to vote for my list but I will not vote unless I am sure my candidates will get their due. If I vote, I am a bad housing board member. Now look what is bound to happen. The blackboard holds the opinions of all the members of the board. The

problem may be settled by voting if the majority of these opinions coincide. Let's calm down and see if it is possible. No, it is not. In fact, we have $N!$ different opinions where N is the number of requests and the probability of coincidence is very small. Consequently, the first thing the members of the board will agree on is that agreement is impossible. By the way, a wise chairman of a housing board would never put a decision to the vote until he is sure it will obtain a unanimous vote. It is perhaps easy to see why. If I disagree and find the arguments given by the other members of the board unconvincing, the session is simply waste of time. I am sure to apply to the next authority up and start the whole thing all over again. As a rule, the decision made by the housing board is to be approved by the higher authority who, for that matter, will send all the papers back to the housing board for their unanimous approval. Alas, such a vote is impossible. Mention should be made of the fact, however, that in some cases voting is not hopeless. If all those attending this lecture elect a chairman from three candidates, we will suceed because we have much fewer alternatives than the number of the people present, and that's when the voting works. In the

> case we are discussing, however, voting is useless. It means the housing board's members must come to an agreement without voting. Then how? First, they may change their minds. Second, we never stop playing with the idea that there is always a possibility to persuade other people to do what we want them to do. To my mind, this idea is naive though I admit that a sober-minded person may be somehow kept from jumping to conclusions or from misjudgement... . It appears that this problem may be formulated in terms of game theory."

Tsetlin knew what he was talking about as he was for many years also a member of the housing board at the Institute of Applied Mathematics of the USSR Academy of Sciences.

Consider a situation in which the conflicts are not as painful as in the housing board. Suppose there is a committee to select examination papers that have been entered for a prize. If the scientific institution organizing the contest does research in more than one field, the department representatives are as a rule convinced that what they do deserves the prize because they know the range of the problems studied, while what the other departments do is a waste of time and effort because they don't quite understand their problems or don't

understand them at all. At the final stage, the work of the contest committee resembles a competition in figure skating where each member of the jury arranges the names of athlets in order and the final result is determined by the sum of placings won by the participants. The members of the jury must be prepared to face a situation in which none of them agrees as to the final placing. Consider the following example which illustrates the situation.

Suppose the contest committee has to study papers written by seven scientists running for a prize. Their names are Johnson, Peterson, Richardson, McHorse, Ratwood, Catwood and Dogwood. Their papers have been submitted by four different departments of a research institute. The contest committee is headed by a chairman who is appointed by the management and consists of four people, one from each department. McHorse is a newcomer at the institute, so he has not become a part of labyrinthine internal politics.

As a result of a long discussion, the contest committee members arranged the names of the candidates for one of the three prizes in order of preference:

1. Johnson	1. Dogwood	1. Johnson
2. Peterson	2. Johnson	2. Ratwood
3. Richarsdon	3. Ratwood	3. Catwood
4. McHorse	4. McHorse	4. McHorse
5. Catwood	5. Catwood	5. Richardson
6. Ratwood	6. Richardson	6. Peterson
7. Dogwood	7. Peterson	7. Dogwood

1. Dogwood	1. Johnson
2. Richardson	2. Catwood
3. Johnson	3. Peterson
4. McHorse	4. McHorse
5. Peterson	5. Richardson
6. Ratwood	6. Dogwood
7. Catwood	7. Ratwood

On the blackboard, the results of the votes were given as

Name	Total placing	Result	Number of votes for a prize
Johnson	8	1st prize	5
Peterson	23	6th place	2
Richardson	21	3rd prize	2
Catwood	22	4th-5th place	2
Ratwood	24	7th place	2
Dogwood	22	4th-5th place	2
McHorse	20	2nd prize	0

The fact that none of the committee members thought McHorse worthy even of the third prize did not prevent him from getting the second pıize. At the same time Dogwood, who was placed first by two members, shares the fourth place with Catwood. It is natural that these results are voted down by at least three votes to two. The chairman suggests the whole procedure be repeated and warns his colleagues that they should be more objective. Now we will see what motives may actually guide the committee members as they attempt to rearrange the list of prize-winners.

The first member is fairly content with the results. Believing that Johnson's pri-

macy is out of the question, he may change his placings of Richardson and Peterson for the better and lower those of their opponents.

Similarly, the second committee member may endeavour to improve the positions of Dogwood and Ratwood. The other three judges act so as to push their own men. The lists produced after the second round of voting are:

1. Peterson
2. Richardson
3. Johnson
4. Ratwood
5. Catwood
6. McHorse
7. Dogwood

1. Dogwood
2. Ratwood
3. Johnson
4. Peterson
5. Catwood
6. McHorse
7. Richardson

1. Ratwood
2. Johnson
3. Catwood
4. Peterson
5. McHorse
6. Richardson
7. Dogwood

1. Dogwood
2. Richardson
3. Johnson
4. Ratwood
5. Peterson
6. McHorse
7. Catwood

1. Catwood
2. Peterson
3. Johnson
4. Ratwood
5. McHorse
6. Dogwood
7. Richardson

The results are again written on the blackboard:

Name	Total placing	Result	Number of votes for a prize
Johnson	14	1st prize	5
Peterson	16	3rd prize	2
Richardson	24	6th place	2
Catwood	21	4th place	2
Ratwood	15	2nd prize	2
Dogwood	22	5th place	2
McHorse	28	7th place	0

The first committee member is quite satisfied by these results and if he changes his vote in the next round he may only spoil Ratwood's results but still can't win a prize for Richardson. It is obvious that the second member could raise Peterson's total placing up to 19 but that wouldn't help Dogwood. Similarly, the third member can't help Catwood at the expense of Peterson and the fifth member can't give a helping hand to Catwood by supressing Ratwood. The results are depressing for the fourth member of the committee, for only one of his favourites is among the prize-winners. However, he is unable to alter the situation either. To everybody's relief, Mr McHorse, the outsider, has lost the game of chance he won in the first round.

The decision obtained in the second round is stable in the sense that none of the judges can change the arrangement on his own without spoiling a result which is for some reason agreeable to him. This suggests that the situation is similar to the Nash equilibrium situation.

On the other hand, these results may be radically altered by coalitions among the committee members. There are different ways to avoid coalitions, one of which is to treat a coalition as an immoral put-up job.

We once tested this multistage evaluating procedure on contest papers submitted by several young scientists. It happened

during a session of the scientific council of the Economics and Mathematics Institute in Leningrad about a decade ago. By the third round only the papers rating eighth and ninth changed places, which made all those present agree that the contest results were fair.

Here the word "fair" means a reasonable compromise between individual preferences and the contradicting preferences of other members of the contest committee. An important feature of this compromise is its stability. Similar procedures could also be used by the housing committee.

4.6. "Stubborn" Automata and Voting

Now consider a formal model for coming to a decentralized agreement. As has become our habit in this narrative, the model

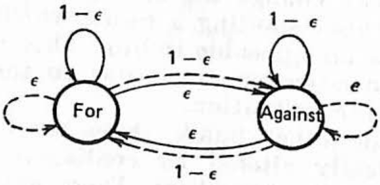

FIG. 4.1

will be described as a collective of automata. The members of this collective are special type of "stubborn" automata, of which the simplest is shown in Fig. 4.1 (as

before, Figs. 4.1 and 4.2 show automata transition after receiving a reward signal in solid arrows and that caused by a fine signal in dashed arrows). We see that a stubborn automaton can be probabilistic. It has two states corresponding to its two types of vote, "pro" and "con". The value of

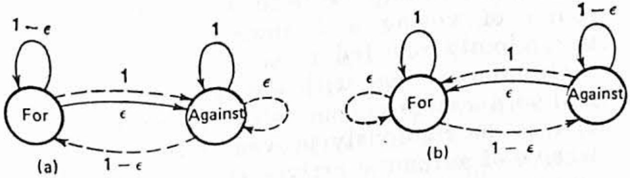

FIG. 4.2

ε characterizes the extent to which an automaton is stubborn. If it incurs a penalty, e.g. it is reprimanded for voting improperly, it might agree and change its opinion with probability $1 - \varepsilon$. In contrast to this, it might insist upon its viewpoint with probability ε. If it is rewarded, the automaton, however stubborn, may still change its mind in the next round with probability ε.

We organize the automata's interaction in accordance with a principle which is very close to the idea of random interaction, i.e. before every round of voting the automata are randomly divided into triples (for convenience we assert that the total number of voting automata is a multiple

of three), and in each triple one equiprobably selected automaton is fined should its state not coincide with the states of the other two automata or rewarded should they all be the same. Further, the selected automaton changes state with a certain degree of stubbornness equal to ε or remains as it is. A new triple is formed prior to every round of voting and another automaton is randomly selected in each triple.

Opinions change with a certain degree of stubbornness to concur with the majority. It may be rigorously proved that this collective of automata arrives at a statistically stable point similar to the Nash point. The number of automata coming to and leaving this point are equal. If during the first round of voting the majority of the automata in the initial collective voted "for", at the equilibrium the absolute majority of automata are sure to vote "for" at the next round. If during the initial round the overwhelming majority vote "against", at the stable point the absolute majority of automata will again vote "against" at the next round. When the numbers of "for" and "against" in the initial condition are approximately equal, the arrival at a stability point will occur only with small values of stubbornness ε.

By reducing ε we can ensure that the automata collective arrives at the stability point (even a one-vote majority in the initial situation will cause enough of

the automata to keep to this opinion to ensure a majority).

Our stubborn automata have been selected so as to estimate the pros and cons symmetrically. Psychological experiments conducted with people have convincingly shown that there is no happy mean in the choice of opinion for an absolute majority of those tested. Some are more inclined to accept positive alternatives and some are more inclined to negative alternatives. It is easy to introduce this asymmetry into an automata collective. The technique is given in Fig. 4.2 in which Fig. 4.2*a* is a diagram of the change in states for a stubborn automaton and Fig. 4.2*b* the diagram for an easy-going automaton. In an automata collective, the behaviours of the symmetrical and asymmetrical automata are identical. Homogeneous associations of both arrive at similar stability points, but with different rates of convergence.

It is natural that an inhomogeneous automata collective is worth discussing. We shall dwell on two types of inhomogeneity. First, we may study an association of symmetrical and asymmetrical automata of both types. Computer-assisted experiments have revealed that such a collective behaves like a homogeneous one. It arrives at both points of stability. Second, we may consider a collective in which automata differ by the degree of stubbornness. The

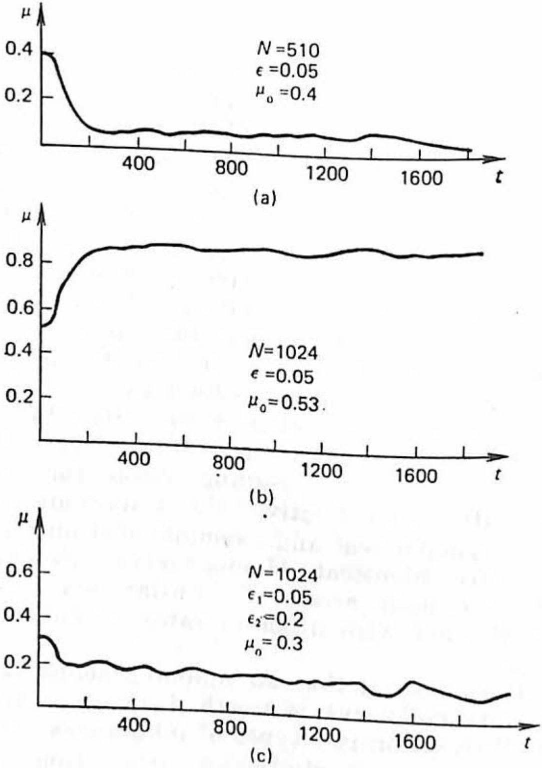

FIG. 4.3

parameter of stubbornness for a certain automaton may be one of the values from the set $\{\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_n\}$. As it follows from the experiment, such a collective still arrives at either a stability point where an

absolute majority votes "for" or at a similar point where it votes "against". Figure 4.3 gives typical curves obtained by the modelling. Along the abscissa we have plotted the number of events t, i.e. the voting rounds, and on the ordinate we have plotted μ, i.e. a segment of N automata in the collective voting "against", where μ_0 is the initial relative part of such automata. Figures 4.3*a* and 4.3*b* are for a homogeneous automata collective having different numbers of automata and different initial preferences within the company of voters. Figure 4.3*c* shows the modelling procedure for an inhomogeneous automata collective in which all the automata fall into two types differing in stubbornness.

It is of interest to note that in a collective where the automata are not mixed up, e.g. the number of likely neighbours each automaton may possess is restricted, the effect observed in Fig. 4.3 is unattainable. This emphasizes the significance of a random interaction procedure in the life of automata collectives.

We have just examined the simplest voting model. At this point we are able to pass over to another model which is reminiscent of the heated discussion in the housing board or the paradoxical triumph after the first round of Mr McHorse, the outsider in the prize contest. As before, our collective will consist of stubborn symmetrical automata. Now each automaton, how-

ever, has K states where K is the number of objects to be arranged in a certain order. Coming back to the example in Sec. 4.5 we see that there $K = 7$ and every stubborn automaton which is to simulate one of the contest committee members must have 7! states. Each state in an automaton corresponds to a certain arrangement of the objects.

We assert that each round of voting is preceded by equiprobable partition of the N automata into $N/2$ pairs; thus random pair interaction comes about. If the collective is composed of an odd number of automata, the remaining automaton votes in the next round according to its old likes and dislikes. For each pair we calculate the value of the difference in opinions or the error variance between the partners in a pair. The value of error variance ρ may be calculated in the following way. Suppose we have two arrangements of preferences (Johnson, Peterson, Richardson, McHorse, Catwood, Ratwood, Dogwood) and (Dogwood, Johnson, Ratwood, McHorse, Catwood, Richardson, Peterson). We hope you remember the eventual prize-winners from the previous section. Johnson, for example, comes first in the first list here and second in the second one. The difference in placing is equal to 1. Peterson comes second in the first scale and seventh in the second arrangement, i.e. the difference is 5. In this manner we go through the two

arrangements and sum the resulting differences. The sum of the differences is the error variance. In our example, $\rho = 1 + 5 + 3 + 0 + 0 + 3 + 6 = 18$.

Consider a collection of N identical automata each with its own preference that is divided into $N/2$ pairs (assuming that N is either even or it represents the number of automata minus one). Then we calculate the error variance for each pair. The total of the values obtained divided by $N/2$ gives us R, which is a mean error variance of the automata collective. It is R that determines the success or failure of the collective interaction. If $R = 0$, then all the automata in a collective with an even number of members have identical opinions.

Suppose in an automata pair there is a randomly selected automaton which will keep to its preferences with probability ε and change them with probability $1 - \varepsilon$. How will its opinions change? The automaton finds which element has the largest contribution to the error variance. In this example, the element is Dogwood with a contribution to the error variance of six. Then the automaton with probability $1 - \varepsilon$ shifts this gentleman several positions so as to reduce the variance. In the simplest case the relocation is by one position. Then, if the change was made by the second automaton in the pair, the arrangement (Dogwood, Johnson, Ratwood, McHorse, Cat-

wood, Richardson, Peterson) is replaced by (Johnson, Dogwood, Ratwood, McHorse, Catwood, Richardson, Peterson). The number of positions by which the element causing the maximum error variance in the arrangements of the pair of automata is moved may be termed the automaton's degree of conformity. The old preferences having been changed (or retained in the most stubborn pairs), another round of voting comes and new random pairs of automata are chosen in the collective.

Computer modelling of this process with different numbers of automata, objects to be arranged, values of a stubbornness parameter, and degree of conformity has shown that this process clearly converges to a single opinion. Typical results of the computer simulation are given in Fig. 4.4. Typical number of modelling events t, i.e. rounds of voting, is plotted along the abscissa and the $\bar{R}$ values are plotted up the ordinate. Figure 4.4*a* illustrates the behaviour of an automata collective with different conformity degrees. It is clear that a higher degree of conformity ω leads to a higher rate of convergence. At the initial moment, the automata's opinions are equally distributed among the acceptable arrangements. In Fig. 4.4*b* you can see the convergence process when the automata's opinions are not initially evenly distributed among the possible arrangements but have a normal bell-shaped dis-

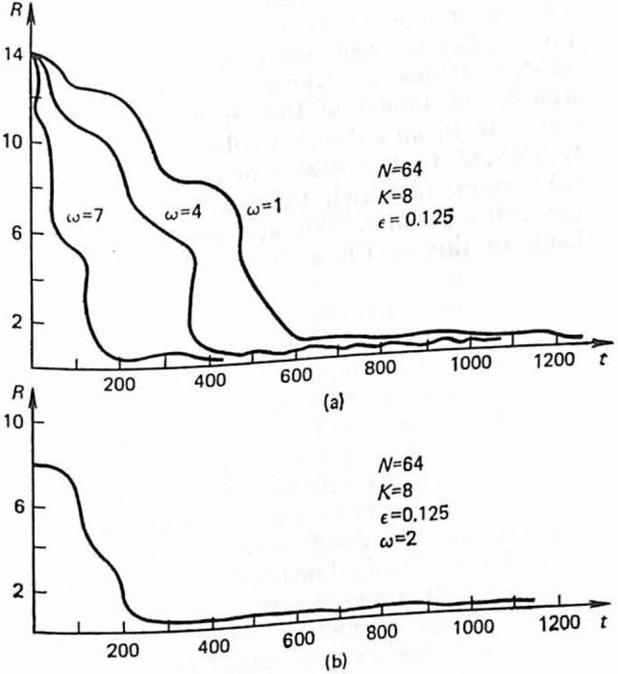

FIG. 4.4

tribution. This is more realistic than the former case because in any group of voters there is usually a coordinated "common opinion" with regard to at least some of the candidates prior to the first round. It follows from Fig. 4.4 that a common opinion existing prior to voting increases the rate of convergence of the opinions.

Mention should be made of the fact that the voting models we have discussed may have different and somewhat unexpected interpretations in different fields. For instance, the model of the simplest pro and con vote in an automata collective is closely related to the design of reliable technical systems in which the main components are self-repairing. We are going to come back to this in Chap. 5.

Chapter 5
Stringless Puppets Make a Show

"The grownup time is now clad in figures."

L. Gongora

5.1. Wait and See Them Fire

It is rather long ago that biologists majoring in culturing tissues, i.e. cultivating living cells from an organism, were studying the synchronization of the moments when the cells split. When a cell splits into two, the two newly-formed cells remain for a while in a state of rest, and then they also split simultaneously. The four new cells also split simultaneously. Several mechanisms may lead to the synchronization of the cell fission. Each cell may either have a precise "internal clock" which determines the intervals between the fissions, or cells may "agree" when to split. However, neither hypothesis has been so far rigorously confirmed experimentally. The first one sounds the more convincing and if a clock does exist, then the mystery of the synchronization is revealed. Synchronization demands that the cells exchange information and a rather rapid signal exchange may be possible if there were, say, a biofield. However, the existence of a biofield can't be regarded as proven experimentally. The speed at which electrochem-

ical processes propagate could hardly ensure the precision with which even quite large cell populations have been observed to split. This is why biologists only concentrated on the two hypotheses to explain the synchronization mechanism, namely, the existence of a biofield and an internal clock. The former hypothesis was, however, problematic, while the latter could only explain synchronous independent fission and suggested considering interaction mechanisms in the case when fission was initiated by some signal coming from outside the cell population and received by a limited number of cells.

A described but unexplained effect occurring in an object under study is discouraging but not fatal. However, when researchers went into the study of self-sustaining models, it was not biologists but engineers and mathematicians who were confronted by the problem of finding mechanisms applicable for causing different parts of a self-sustaining machine to be turned on simultaneously. We must admit that this does not seem an insoluble problem to an engineer; it is enough to introduce a sufficiently precise clock common for the whole system which can inform all the components of the system of the current time. In technology, this method of temporal interaction between the individual behaviour in the system is called *synchronization*. The idea of synchronization based

on a common clock, however, seemed to be different from the biological case, and so the American scientist J. Myhill went on to pose his famous "firing squad synchronization problem". Myhill's problem is formulated as follows.

Consider a chain of soldiers (Fig. 5.1), each soldier being able to exchange in-

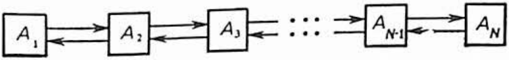

FIG. 5.1

formation only with his two immediate neighbours. The chain is finite, so the two soldiers at the ends of the chain have only one neighbour each. One of the end soldiers is given a command after which the whole squad must "come to an agreement" and fire their weapons simultaneously. The question is whether there exists a local algorithm of behaviour which enables the soldiers to do this if the number of words the soldiers may use to convey the information and the memory capacity of each soldier are finite and independent of the number of soldiers in the chain.

The assumption that there may exist a solution of Myhill's problem suggests that it may be possible to synchronize an infinitely long chain of limited complexity objects by organizing their interactions

with the signal exchange going on at an infinitely low rate.

Myhill's problem was first solved by J. McCarthy and M. Minsky and then in 1962 E. Goto found a solution where the synchronization took place within a minimal time $T = 2N - 2$, where N is the number of soldiers. The behaviour algorithm for each soldier was represented by an end automaton with several thousand inner states. In 1965, the Soviet researcher V. Levenshtein published a brilliant solution with a 9-state automaton. Further research helped reduce the number of states to eight.

Although the solution of Myhill's problem gave an answer to a significant methodological question and demonstrated a number of effects which will be considered in the coming section, many engineers treated the obtained structures sceptically asking what all the fuss was for if you can pull a wire connecting all the soldiers and switch on a lamp for all of them thus giving them a signal to fire.

Until recently, it was not easy to argue with such sceptics. The principle of external synchronization solved a great many technical problems perfectly. As time went on, however, it appeared more and more difficult to rely on common synchronization when dealing with very complex systems. The advent of submicron integrated circuits has aggravated the situation. If

a transistor junction in a crystal is shorter than a micron, the delays in the connecting conductors become significantly longer than the transistor's switchover time, so we are again in a situation in which the information exchange between the objects becomes relatively slow. A new term, i.e. *equichronous zone* has appeared. It means a zone in a crystal where it may be assumed that time flows equally. It should be noted that the synchronization of processes occurring in different equichronous zones requires the organization of a special interaction between the zones.

There are a variety of approaches to this problem and we now turn to considering a solution connected with Myhill's problem.

5.2. Have Them Fire All at Once

A formal statement of Myhill's problem in terms of the automata simulating the soldiers can be given as follows.

Consider a chain of N automata each having n inner states. The state of each automaton at each successive moment in time depends on its own state at the present moment and those of its right-hand and left-hand neighbours. At the initial moment all the automata are in an initial state S_0. At this moment, one of the end automata in the chain receives a signal from the environment which changes its

initial state. The question is whether there exists a structure for each automaton (or rules which govern the changes in state) such that initiating the end automaton in the chain is eventually followed by the simultaneous transfer of all the automata

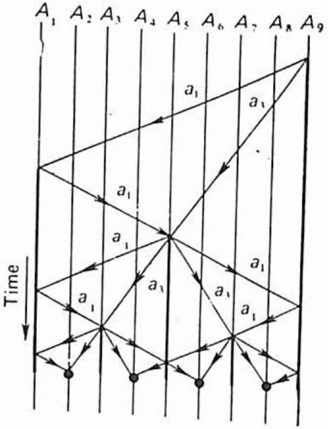

FIG. 5.2

to the same state S and none of them enters this state before this moment. Moreover, the complexity of each automaton n is independent of the length of the chain.

Consider a possible algorithm for regulating the automata interaction which provides a solution to this problem. Figure 5.2 is the time diagram for a 9-automata chain. Each automaton is given by a vertical line. A thin line represents the initial state S_0

of the automaton. The initiating signal comes to the end automaton A_9 of the chain and transfers it into a state called the state of readiness S_1. After it has changed to this state, the automaton produces two signals a_1 and a_3, which start propagating down the chain with the velocities of 1 and 1/3 respectively. The propagation of a signal down the chain with a velocity of 1 means that an automaton receiving signal a_1 from the right transmits it in the same direction in the next time unit and if a signal has a velocity of 1/3, it delays signal a_3 for three time units.

Having reached the opposite end of the chain, signal a_1 changes the state of the second end automaton into the state of readiness S_1, is reflected from the end of the chain, and begins to propagate in the opposite direction, i.e. from left to right. It is not hard to see that the reflected signal a_1 and signal a_3 are bound to meet at the centre of the chain. Automaton A_5 which is located at the meeting point (or the two central automata if there is an even number of them in the chain) enters the state of readiness S_1, which is denoted in Fig. 5.2 by a thick line.

Automaton A_5, which is now in state S_1, sends out two pairs of signals a_1 and a_3 in both directions; note that signal a_1 is reflected from the first automaton in state S_1 it encounters. As a result, at the points where the reflected signals meet signal a_3 (A_3 and

A_7), the automaton changes to state S_1 and another two sets of the a_1 and a_3 signals are generated in both directions.

Thus, in every cycle the interval between each pair of automata in state S_1 is divided into two, and in the centre of the interval the automaton also enters state S_1. Note that prior to the last division each automaton has at least one neighbour which is not prepared for synchronization, i.e. is in a state other than S_1. It is assumed that an automaton enters the synchronized state S if it and its two immediate neighbours are in state S_1. As you see, the synchronization problem is soluble and the rules governing each automaton's behaviour do not depend on the length of the chain. The time needed to divide the interval into two equal parts amounts to 3/2 of its length; thus the total time needed for synchronization, give or take constant interval depending on whether the number of automata is odd or even, is equal to three lengths of the chain.

The synchronization time may be reduced if an automaton in a state of readiness also sends out a signal a_7 which propagates with a velocity of 1/7, and a signal a_{15} which spreads with a velocity of 1/15 (cf. Fig. 5.3). The reasons for the faster synchronization process in a situation like this become obvious after a comparative analysis of Figs. 5.2 and 5.3. Although in this case, too, the local behaviour algorithms

do not depend on chain length, the automata tend to grow more complex.

A researcher who wants to teach an individual automaton to delay a signal for T time units must make sure that it has T or more inner states. If he wants this automaton to delay one signal for T_1 time units and another signal for T_2 time units, then

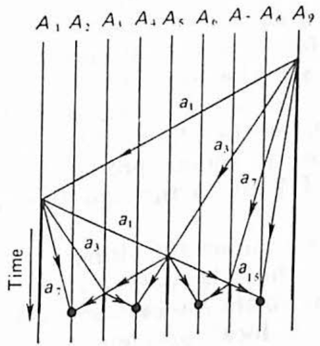

FIG. 5.3

the number of inner states must be at least T_1T_2. Thus the introduction of additional signals to speed up the synchronization requires a far more complex automaton. Moreover, now the necessary number of signals and, consequently, the automata complexity become dependent on the chain length.

In the previous section we mentioned Levenshtein's brilliant solution of the synchronization problem. His solution was

very elegant for he managed to find a method of organizing the automata interaction in such a way that in the course of interaction between neighbouring automata the signals are delayed for a number of time units equal to 1, 3, 7, 15, . . ., $(2^k - 1)$, where k depends on chain length but is independent of the number of each automaton's inner states, the automata having only nine inner states. At present, a solution for 8-state automata is known. In this case the synchronization time is reduced to its minimal value that is twice the length of the chain, i.e. the time necessary for a signal propagating with a velocity of 1 to go through the chain and back.

The most important thing in this solution from the viewpoint of methodology is that due to its interaction each automaton solves a local problem which is more complex than the one it could solve on its own. Moreover, as the number of automata in the chain grows, each one appears to be able to cope with more and more complex local problems by relying on an interaction with its two immediate neighbours.

Myhill's problem gives rise to a whole class of interesting problems. It includes the synchronization problem in which the starting signal is applied to an arbitrarily selected automaton. In this case the underlying automaton has ten inner states.

Another question brought about by the study of models of this type is whether it is possible to synchronize the automata with an unknown delay between them and with the automata complexity being independent of the delay. The last requirement rules out a simple algorithm for obtaining a time interval between the emission of a signal and the reception of its echo. Similarly, this algorithm can't be used if the delay changes in time. Suppose we want to synchronize events hapenning on the Earth with those aboard a spacecraft moving rapidly away from it. In an exotic situation like this, we are to synchronize the system by arranging for the components to send synchronization signals to each other at the same moment. The question is whether there exists an algorithm for reaching a synchronized state through the local behaviours of a ground flight control system and spacecraft systems. In terms of the automata models discussed, the solution appears rather simple and may be obtained by means of 12-state automata.

A synchronization algorithm which will work to one time unit is relatively simple. The automaton initiating the synchronization sends three signals a_1, a_2, and a_3 at three consecutive moments in time. The receiver sends these signals back delaying signal a_1 by one time unit, signal a_2 by two time units and signal a_3 by three time units. The simultaneous reception of signals a_1

and a_3 by one of those taking part in the exchange means a synchronized state; and it starts sending the synchronized signal s onto the communication line. When signals a_1 and a_3 are received by one automaton, the other one receives signal a_2 and turns it back with a delay of two time units. When a synchronized state is achieved, signal a_2 is no different from signal s and if the latter

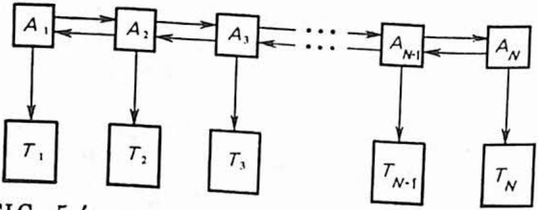

FIG. 5.4

comes back with a delay of two time units, then the exchange of the s signals maintains the synchronization. To avoid the effect of probable errors in the local clocks showing the local time on the Earth and aboard the spacecraft, the a_1 and a_2 signals must also continue to be exchanged after reaching the synchronized state.

Now let us turn to a consideration of a situation in which the soldiers in Myhill's firing squad do not fire their rifles but switch on some devices (Fig. 5.4), each of which has its own latent time, i.e. a time delay between pressing the "ON" button and the moment it starts operating. Thus,

it takes a minute from the moment a radar emitter is turned on to the moment when the radio valves are heated up, while it takes five minutes for a petrol engine to warm up. Each "rifleman" knows the latent time of his own device but only his. The question is whether there exist local rules for automata behaviours the complexity (the number of states) of which does not depend either on the length of the chain or on the latent times of the other "riflemen". In other words, when the command to fire is given to one of the soldiers, they all must press their "ON" buttons so that their devices start working simultaneously.

It appears that this problem has a solution based on the principles of Myhill's initial problem and the underlying automaton is a chain of T_k automata handling Myhill's problem, where T_k is the local latent time. This arrangement is supplemented by logic which produces a start signal according to the states of the automata in the above chain.

Since we are considering algorithms of local limited interaction, it is natural that we are interested in the effect the number of neighbours has on the problem's solution. A study of this problem has revealed that synchronization in a segment is nearly as difficult as in an arbitrary graph. On the other hand, it is natural to ask whether synchronization is possible when only one neighbour is available. This

question has a positive answer. Any problem soluble in a segment with two neighbours can be handled in a circle with only one neighbour.

5.3. Marching and Wandering Automata

There is a trick in the problems considered. The models solve a synchronization problem by utilizing an inner time unit, which comes about like a godsend. It is only when considering the synchronization of a system with a delay of unknown length in the communication lines that we found a probable mechanism for coordinating the lengths of the inner time units.

It is impossible to disregard the notion of an inner time unit as long as we deal with the finite-state automaton model used to describe the local behaviour algorithms since the model takes no account of real physical time. The moment we start regarding an automaton as a real device, it becomes a dynamic system in which the transition processes differ from the procedure of the change in state that occurs in a finite-state automaton.

We may overcome these difficulties through the use of the theory of integrated self-synchronizing circuits, which ensure the invariance of the automata's behaviours with respect to the length of the inner time unit. An account of the ideas and results of this theory, however, lies far beyond

the scope of this book. However, we can consider a model in which synchronization is reached without relying on a rigid time unit.

In Fig. 5.4, the "riflemen" make up a line, which is a rigid organization. The question now is whether it is possible to synchronize a "crowd". It is assumed that the automata wander within a certain limited space and once in a while run into each other. Unlike random pair interaction models, we do not assume that these encounters happen at the same time, thus avoiding the need for synchronization. Our model or, to be more exact, its formal study demands that the automata crowd be well mixed, i.e. that ideally all the encounters be equiprobable. Mixing brought about by processes of the Brownian motion type slightly distorts the results. During an encounter, two automata interact and part. The question arises as to whether there exist interaction rules which ensure the synchronization of all the automata and which are independent of the number of automata.

First we will specify what we mean by synchronization in this case. Initially all the automata are in some state S_0 and when two automata come to contact they remain in this state. The outside signal starting the synchronization process is applied to one of the automata and changes its state. It is natural that we can't expect all the automata to go into a synchronized state

after a fixed period of time because in random interaction there is always a nonzero probability that during a finite time the information concerning the initiation of the process will not go beyond a limited number of the automata. Therefore, in random interaction we can only speak of the probability with which a considerable proportion of the automata will enter a synchronized state either simultaneously or within a relatively short period. The length of this period is dependent on the frequency of encounters. Since the automata change state due to interactions, all the synchronized automata must have been involved in an interaction within this time interval.

We will estimate the quality or accuracy of the synchronization by the mathematical expectation of the number of automata which are in the synchronized state. We assume that there is a certain number ε $(0 < \varepsilon < 1)$ and that we can determine a time interval t_ε during which a part of the automata exceeding $1 - \varepsilon$ must interact. We will call an automata collection ε-synchronizable if a randomly selected automaton has been initiated at t_0 and there exists a moment t_c such that prior to this moment the mathematical expectation of the synchronized portion of the automata is less than ε, while the mathematical expectation of the synchronized portion of the automata after $t_c + t_\varepsilon$ exceeds $(1 - \varepsilon)$.

Now let us turn to a consideration of a local interaction algorithm which ensures the solution of this synchronization problem. It is assumed that an automaton is in state $(k-1)$, the state numbered 0 being the initial state and state numbered k being the synchronized state. The interaction procedure is governed by the following rules:

1. If both interacting automata are in state 0, they remain in this state.

2. An outside initiating signal transfers an automaton from state 0 to state 1.

3. If at least one of the interacting automata is in a state other than 0, both automata change to a new state whose number is the lesser of the two state numbers of the interacting automata plus one.

We now turn to an analysis of the synchronization procedure. The first automaton, which is initiated by the outside signal, goes into state 1 and transfers all the other automata, which were in state 0 to state 1. The number of transitions to state 1 in the automata collection snowballs spontaneously. An automaton in state 1 will go into state 2 if it meets another automaton in state 1; the probability of this is equal to a proportion of automata in state 1. Moreover, if an automaton in state 2 encounters an automaton in state 0, it will return to state 1. Hence, if the rate of state 0-state 1 transitions linearly rises with the proportion of automata in state 1,

then, roughly speaking, the rate of state 1-state 2 transitions grows as the square root of this proportion. In general, if the state numbers are somehow distributed among the automata, the rate at which the proportion of automata with smaller state numbers decreases is much higher than the rate at which the proportion of automata with larger state numbers increases. We can expect the spread to grow as this state-number distribution goes on. In other words, the state-number distribution of the automata will "concentrate" around the number of the current average state.

The validity of this assumption is proved both by mathematical analysis and by the computer simulation of the behaviour of such an automata collection. During the simulation, one time unit was the time it took all automata to interact once. Thus, for example, in a collection of 1024 automata each having 15 inner states, all the automata entered the synchronized state in 400 experiments. It is crucial here that there is a mathematical proof that for a given ε the number of inner states required for each automaton is independent of the total number of automata in a collection. The fact was confirmed for large automata collections and small ε. This however, does not make the result obtained less meaningful.

These results seem amazing. As the size

of the automata collection grows, we could expect the problem of synchronizing its behaviour to become more complex. There exist, however, rather simple rules of random interaction which help us synchronize the response of an automata collection to information supplied to one automaton. This effect is produced by processes in which those lagging behind are accelerated and catch with the rest, while those moving too fast are slow and wait for the rest to catch up with them. These processes may be defined as *phase synchronization*. It is not improbable that we can come across similar processes during evolution. Random changes occurring in separate individuals can spread throughout the population due to random pair interaction and can be realized in a synchronized manner when the situation is favourable.

5.4. Praise Be to Homogeneous Structures

The reader must have noticed our weakness for homogeneous structures interacting in the same fashion. This is indeed the case. It is with surprise and admiration that we discover the abilities in these simple structures. A unidimensional homogeneous structure has just demonstrated its synchronizability without any global synchronizing signal. In Chap. 3, we watched irrigation pumps which could turn themselves on in the manner we wanted them

to. Now we are going to discuss some more examples from the wonderland inhabited by homogeneous structures. It is noteworthy that these structures, which have been known to the biologist, now excite great curiosity in engineers. Indeed, they are much easier to manufacture, replace, and control. Identical homogeneous structures are something an engineer can only dream of. That's why in recent years and especially after the advent of microelectronics homogeneous structures have become a matter of great interest.

We believe that the scene described below is well-known to all readers. A heated debate of a controversial issue is in full swing in a conference-hall. The chairman puts the motion to the vote, asking those in favour to raise their hands. Counting votes is tedious and time consuming. Sometimes the chairman can't do the job alone and needs assistance. Then he asks those against to raise their hands and the counting procedure starts all over again. It is best if one of the two opinions has only a few supporters. In this case there is no need to count the votes at all for the result of the vote is obvious. If the motion has as many supporters as it has opponents, however, passions among the audience may run high. The voters may demand a recount since those who have lost always think that the count was inaccurate. Can this awkward situation be avoided? Is it possible

to count the votes and determine the result by some sort of automatic device? Can we build a "voting machine" to produce the result whatever the number of voters? All these questions may be answered positively. We can handle these problems by using unidimensional and two-dimensional homogeneous structures. To prove this statement we now turn to very simple examples.

First consider the most elementary way of voting in which a motion is either supported or opposed and is adopted by a simple majority vote. It is obvious that in any calculation every "for-against" pair may be crossed out. After we have crossed out all such pairs, we need only deal with those who voted "for" or those who voted "against", depending on which of the two is in the vote majority. There is no need to count the votes, for in this case a one-vote majority is enough. The procedure of crossing out these pairs of opponents may be easily realized in a unidimensional homogeneous structure. Consider the automata collection given in Fig. 5.5. Each automaton has two states, denoted a and b. State a means that the result of voting has not yet been crossed out and state b that it has. Signals x_i come to automata inputs from the voters. If $x_i = 1$, the ith voter has voted "against" and if $x_i = 0$, he has voted "for". If the voting procedure takes place in a specially equipped room, then a signal

from a voter comes to our device by his pressing an appropriate button on a panel.

Signals sent through a horizontal bus between the automata are internal (operat-

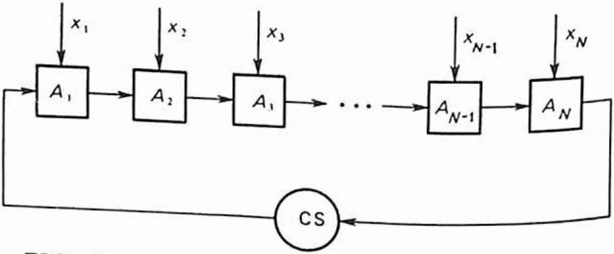

FIG. 5.5

ing) with respect to our device. The operation of the automata in the chain is given in Table 5.1.

Table 5.1

Operating signal	Input signal from the ith voter			
	$x_i = 0$		$x_i = 1$	
	state of automaton A_i		state of automaton A_i	
	a	b	a	b
s_0	s_1, b	s_0, b	s_2, b	s_0, b
s_1	s_1, a	s_1, b	s_3, b	s_1, b
s_2	s_3, b	s_2, b	s_2, a	s_2, b
s_3	s_3, a	s_3, b	s_3, a	s_3, b

Here are some clues to the table. At some moment during the operation, automaton A_i receives an input signal x_i, which is either 0 or 1, and one of the four possible operating signals. A_i itself is in one of the two possible states, a or b. The values of input and operating signals and the automaton's state determine the current situation, which is written as a pair of symbols: value of the operating signal produced by automation A_i in the current situation and the new state which it is entering.

At an early stage during the operation, the control system (CS) which is labelled in Fig. 5.5 by a circle, transmits the signal s_0 to the input of the left-most automaton in the chain. During the initial time unit of operation, all the automata are in the noncrossed-out state a. When the first automaton which is in state a (in the initial time unit this is the left-most automaton in the chain) is encountered, the signal s_0 transfers it to the crossed-out state b. If the signal sent to the automaton's external input is a "for" signal then the s_1 signal will travel farthest to the right end of the chain. If $x_i = 1$, i.e. an "against" signal, is at the automaton's input, the s_2 signal will travel to the right.

The s_1 and s_2 signals are "looking" for the first automaton to make up a pair to be crossed out with an automaton they have discovered. If the s_1 or s_2 signals manage to

find an appropriate automaton on their way to the right end of the chain, it is transferred into the crossed-out state and the operating signal becomes s_3. This signal indicates that one pair of the multitude of voters is crossed. The s_3 signal "jumps" through the rest automata of the chain without being changed on the way. Its arrival at the control system indicates the end of one crossing-out event. This arrival causes the CS to send a new s_0 signal to the automata chain. Crossing-out is over when the signals s_1 or s_2 can't "pair off" so as to be crossed out. In this case instead of an s_3 signal, the s_1 or s_2 signals will be input to the CS. The arrival of an s_1 signal means that a majority voted for the motion and the arrival of an s_2 signal means that their opponents won. A special case is a stalemate when the numbers of voters on either side are equal. Then the control system sends the s_0 signal to the input of the left-most automaton of the chain. This signal "jumps" through all the crossed-out automata intact and arrives at the input of the CS. Thus a stalemate is indicated.

Note that the value of N has been neither specified nor taken into account. However many voters there are, the vote counter will function properly.

So we have considered the simplest case of voting. In more complex cases, too, the problem of vote counting is soluble on a unidimensional automata chain in which

the required complexity of each automaton does not depend on the total number of votes but only on the type of voting. Thus, for example, in a two-thirds' majority vote, the automata complexity is not subject to change, while the number of operating signals grows up to six. One scan by an s_0 signal causes an automata triple to be crossed out if two automata have a "for" signal and one an "against" signal. We could as well consider a system in which one of K candidates must be elected by a majority vote. This problem demands a 4-state automata and eight operating signals. A consecutive voting system (for example, you may choose either B_1 or the pair (B_2, B_3) from the candidates B_1, B_2 and B_3; if B_1 is chosen, then the voting is over; if the (B_2, B_3) pair is chosen, then B_2 and B_3 compete in the second round) requires a two-dimensional homogeneous structure. The same may be said about a voting procedure in which each candidate running for a prize gets a certain number of points.

Voting procedures invented by man are many and diverse. What is noteworthy is that all of them (at least all those known today) may be simulated by means of homogeneous structures whose structural complexity does not depend on the number of voters.

At present there are many applications of homogeneous structures in a variety of fields. Their use, however, is somehow limit-

ed owing to the fact that they realize familiar things in somewhat unusual fashions. Take, for example, multiplication. Since our school days, we have multiplied numbers in column believing that this method is least time consuming. In fact, this is not quite so. The breakthroughs that have been made in computers have made the decimal number system give way to the binary system. In the binary system, multiplication is far simpler. A combination of consecutive adding and shifting by one position produces the same effect as multiplication in a column, without the need of a multiplication table. Below we show you how 12 is multiplied by 14 in the binary system:

$$\begin{array}{r} \times\; 1100 \\ \underline{1110} \\ 1100 \\ +\;1100 \\ 1100 \\ \underline{0000} \\ 10101000\,. \end{array}$$

Here we multiplied beginning from a most significant digit downwards. We could as well multiply by starting from the least significant digit and going up. The result would have been the same. As in the decimal system the product is 168.

Figure 5.6 presents a homogeneous matrix made up of three columns, each square of which contains single type of automaton

which has three states, D_0, D_1 and D. D_0 and D_1 are called operating states. If an automaton is in either of these two states, it stores in this square the value of a binary number, viz. 0 and 1 respectively. D means an "off" (idle) state. Each automaton is connected with all its neighbours.

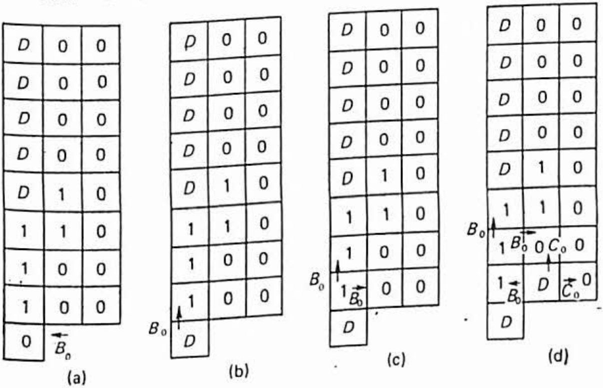

FIG. 5.6

We will designate these inputs and outputs by the letters b, u, l, r (bottom, upper, left, right). They will be used to transmit five types of signals: idle (useless) and B_0, B_1, C_0, C_1. We have intentionally given no letter to the idle signal. The signals B_0 and B_1 inform you of the digits in the corresponding positions in the multiplier. The C_0 and C_1 signals generated by the automata correspond to the transmission of the signals 0 and 1 to a neighbour.

Table 5.2

Number of command	Current time unit			
	input			state of automaton
	r	l	b	
1			C_k	D
2			C_k	D_p
3	B_0			D_p
4			B_k	D_p
5			B_k	D
6		B_1		D_p
7		B_0		D_p
8		B_1	C_k	D_p
9		B_0	C_k	D_p
10	B_0			D
11		C_k	$^*C_p$	D_i
		C_k	$^*C_p$	D_i
		C_k	$^*C_p$	D_i
		C_k	$^*C_p$	D_i

The number of lines in the matrix varies with the number of digits to be multiplied together. With n positions, the number of lines in the first column is equal to $2n + 1$ and in the second and third columns $2n$. The matrix in Fig. 5.6 is designed to multiply numbers with four binary digits.

All the automata in the matrix function in the same manner. This principle stems from the homogeneity of the environment. The automata operation is illustrated in Table 5.2.

Here the indices k, p, i can take the val-

Coming time unit				Transition condition
input			state of automaton	
r	u	l		
			D_k	
	C_k		D_p	
	B_p		D	
B_k	B_k		D_p	
B_k	B_k		D	
C_p	C_p	B_0	D	
C_0	C_p	B_0	D	
C_p	C_p		D_k	
C_0	C_p		D_k	
			D	
	C_1		D_1	$k+p+i=3$
	C_1		D_0	$k+p+i=2$
	C_0		D_1	$k+p+i=1$
	C_0		D_0	$k+p+i=0$

ues 0 or 1. Only those places in the table are filled which correspond to the combinations of signals and states found in a multiplication process. An asterisk means that there may be no such signals.

How does multiplication take place? Figure 5.6*a* shows how the matrix is initially filled. The multiplier appears in the left-hand column so that the least-significant digit is at the bottom. The multiplicand is written in the next column. The right-hand column is used to write the product. The idea is as simple as that: if the

next digit of the multiplier is 0, then the multiplicand must be shifted one digit up. If the next digit of the multiplier is 1, then, before the multiplier is shifted, it must be added to the current running in the right column of the matrix.

The process starts from the generation of the signal B_0 from the right side to the lower digit of the multiplier. This signal "reads" the digit's value. If it is followed by the appearance of a D state, the digit is read (cf. Fig. 5.6*b*). Further, the signal B_0 travels up the left column and at each level sends a signal B_0 to the right. These signals shift the multiplicand with no addition to the running total of the matrix's right-hand column. In the example considered this happens because the multiplier digit in the initial cycle was equal to 0. Figures 5.6*c* and *d* give two more time units in the operation of the homogeneous environment. In Fig. 5.6*d*, the appearance of the B_0 signal coming to the left-hand column indicates the beginning of the operation with the multiplier's next digit. Thus, during the fourth time unit the multiplication microcycle is over.

A general view of the signal spreading through a homogeneous matrix in a binary digit multiplication operation in column is given in Fig. 5.7. Here a double arrow shows the B_k signals; an ordinary arrow indicates the C_k signals; an open circle indicates an automaton's transition during this unit to

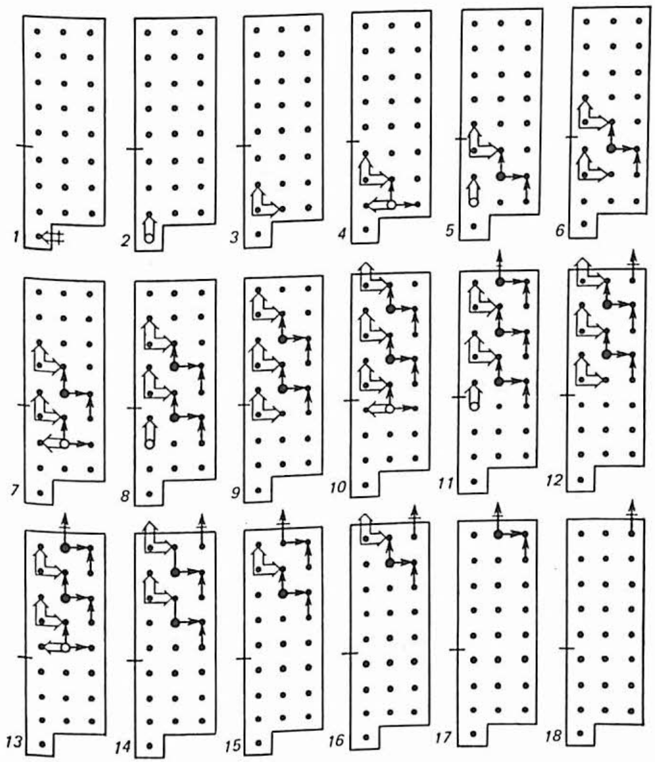

FIG. 5.7

the D state and the transmission upwards of the value of its state; and a bold circle shows an automaton's transition to the D_p state in accordance with the value of the signal arriving at that moment from the bottom. A crossed arrow means that either

a B_0 or a C_0 signal is generated. It follows from Fig. 5.7 that a multiplication cycle of two numbers with four binary digits is completed in 18 time units. The result of the multiplication is in the states of the right-hand column automata. Signals which in Fig. 5.7 look as if they were going beyond the matrix are lost and have no effect on its further operation. The reader can complete multiplying 12 by 14 using a homogeneous matrix consulting Table 5.2 and Fig. 5.7.

The total number of time units required to multiply two numbers with n binary digits using the method described is $8n + 2$. However, if we avoid traditional logic altogether and instead pass over to more "exotic" multiplication techniques, a homogeneous matrix allows us to get the product in $4n - 1$ time units. Here the automata which make up the matrix are even simpler than in column multiplication. True, the multiplication of two n-bit numbers requires $n^2 + 1$ automata, whereas in the case above only $6n + 1$ automata would be needed.

Besides multiplication, homogeneous matrices are helpful in division too. A division method using a two-column homogeneous matrix demands automata of the same complexity as does the multiplication method. The number of automata needed is $2n + 3$. Homogeneous environments appear very promising in a variety of fields.

Thus, they prove indispensable when processing visual information in devices similar to photoelectric matrices. In this case homogeneous automata structures allow us to determine the outline of the picture, find a vertex of a cone, a crossing point, or a distance between pictures in a matrix. The recent breakthrough in robot application has further stimulated this line of research.

An important thing here is that homogeneous structures and groups of simple devices can easily cope with problems which we traditionally assign to consecutive or centralized methods. The difficulties we encounter when passing over to parallel or decentralized methods lie in that techniques and algorithms realized in homogeneous structures seem unusual.

This praise of homogeneity does not contradict our previous arguments in favour of inhomogeneity. In Chap. 3, we sought to demonstrate the new qualities introduced by inhomogeneity to the behaviour of an automaton collective when solving a common problem. It has not yet been proved that homogeneous collectives and structures can efficiently cope with all the problems we wish our machines to solve. No researcher, however, has so far proved the opposite!

5.5. Why Yoga Is Not Our Way?

This was the title of a scientific report delivered by M. Bongard, a well-known Soviet expert in cybernetics, at a conference devot-

ed to collective behaviour models. In this report, he spoke about the great harm that may be caused to a living organism by excessive centralization. Increasing centralization will make an organism waste a growing amount of resources to process information leaving less and less time for search and adaptation activity. As an example, he mentioned advanced yogis whose practice, among other things, includes "raising the level of conscience" to control physiological processes which generally occur in humans at the level of autonomous or semi-autonomous control systems. For example, they can speed up or slow down their cardiac rhythms, cause their stomachs to contract or relax, and change their body temperatures. What effect do these abilities produce? If they were to neutralize all the automatic systems, a yogi would be compelled to use all his time and brain resources to avoid a malfunction which may threaten his life. He would then have no time left for meditation or contemplation. It is certain that Hindu yogis never find themselves in such a dangerous position. They don't neutralize the automatic reflexes entirely and only occasionally interfere in their operation. The point is they have another goal. By trying to unveil the secret of control over the autonomous processes occurring at a subconscious level, they seek to comprehend the laws which govern their hu-

man bodies. However, Bongard's analogy is both vivid and instructive.

We have spoken much about parallel processes and methods of their interaction. In a human body this interaction is much more complex. The nature of the phenomenon is, however, the same. The processes occur almost autonomously and are synchronized in time by occasional periodic or specific signals generated in particular situations.

We have to admit that decentralization, in which subsystems operate practically independently of any control centre, has a significant drawback which we have not mentioned yet. The reader must have noticed this drawback himself since in many of our models working in rapidly changing environments it was obvious. We mean a longer period of adaptation which is a penalty for decentralization. A decentralized system reacts more slowly to any change in the environment than does a system in which commands are generated by a centre having advanced information about the coming changes in the properties of the environment. It is probably for this reason that living organisms and, particularly, man has two levels of control, centralized and decentralized. These levels, however, do not overlap.

Decentralized control completely dominates the situation when the environment is favourable and is not subject to significant

change. Our separate subsystems function on their own and have almost no interaction with each other. Then comes a rapid change in the environment which is fraught with dangerous consequences. It is vital that all the subsystems be put in a "state of alert" as soon as possible. This is the time for a centralized control to function and put a human body into a state of stress. The main feature of this reaction is that it is not specific. Whatever the danger an organism is facing, this reaction causes all the subsystems of the body to interact with each other. Blood starts receiving hormones which improve adaptability, the body is warned of the coming loss of energy, muscles are better fed, etc. After that, either we adapt or the stress situation comes to an end. In the worst case, a human body remains in a state of "immediate standby" for so long that it becomes exhausted or even collapses.

Thus here we see an intriguing distribution of functions between the decentralized and centralized parts of a control system. In slowly changing or constant environments, a decentralized control system can cope with the task of adaptation and achieving the body's global goals, while a central controller is turned on whenever there is a need to handle a dangerous contingency.

Experts in the control of integral robots (which, unlike highly-specialized robots, must be able to operate within a wide range

of environments which cannot be accurately described) are now at a loss. On the one hand, it is obvious that a robot must have several subsystems which must function autonomously or nearly autonomously after receiving signals from a central controller (for example, "eye" and "hand" subsystems which allow a robot to find a desired object, grasp it, and perform some operation with it, which act parallel to each other and autonomously, while interacting in accordance with synchronizing signals). On the other hand, there is a need for non-specific global commands generated by a central controller and capable of providing the robot's proper behaviour. It is extremely difficult to formulate the general laws for such behaviour. Let us recollect the three general laws of robotics suggested by Isaac Azimov. These laws are given in order of priority. The first and most important law states that a robot must under no circumstances do any harm to a human being. This is a law of total ban. It is not hard to imagine what should be done if there is a danger that this law should be violated. Azimov's second law states that a robot must obey a human being unless this runs counter to the first law. The third law of robotics says a robot must protect itself unless it runs counter to the first two laws. The last two laws, however, can't be non-specific with regard to signals sent to the robot's subsystems. They must be

more specific about the type of orders which can be given to a robot and its self-protection techniques.

The non-specific signals to be generated by the centralized part of a control system were, for many of the automaton and non-automaton behaviour models in this book, presented as being influences of the environment on the subsystems. Mechanisms such as a common fund procedure of random pair interactions had general control functions. We mentioned in Sec. 4.4 that a collective's goal may not only be the achievement of an appropriate or optimal behaviour in the environment but also a search for regulating effects which permit the subsystems to come to a coordinated regime.

To show you how harmful it is to try to "pull" specific functions into the centralized part of a control system, we would like to end this chapter with an anecdote which really happened at an international conference on artificial intelligence and robotics. When one researcher lamented that it was rather difficult to invent a limited number of non-specific laws to generate expedient behaviour for integral robots, an important US Navy representative remarked that he did not think the job was that hard. To prove it, he gave the following example. When a novice arrives aboard a ship, he can't adapt to the environment right away and is bound to make many mistakes. While trying to be helpful, he

may damage the ship. The experienced crew thus have to waste a great deal of effort to watch him and save from trouble. However, it is easy to avoid the problem. It is enough for a novice to remember one non-specific law for the whole period of his adaptation aboard a ship: "If it moves, salute it; if it doesn't, paint it".

This is a joke but it describes our problem quite accurately. Today, however, we do not know much about how to construct this sort of regulating procedure in decentralized control systems.

Chapter 6

Dialectics of the Simple and the Complex

"You should see the difference between roads you choose and roads you are chosen by".

Felix Krivin

6.1. Synthesogenesis and Integration of Efforts

"... The strictly symmetrical tripartite structures resembled the letter Y. Three wings were anchored in a central thicknening, each wing tapering to a point at its extremity. They looked coal black under direct illumination; but reflected light made them glisten bluish and olive green, not unlike the abdomens of certain terrestrial insects which are composed of tiny surfaces like the multifaceted rose-cut of a diamond. Their interior structure was always the same when examined under a microscope. These minuscule elements, one-hundredth the size of a small grain of sand, formed a kind of autonomous nervous system with a number of independent fibres.

"The smaller section, forming the arms of the letter Y, constituted a steering system controlling the "insect's" locomotion. The micro-crystalline structure of the arms provided a

type of universal accumulator and at the same time an energy transformer. Depending on the manner in which the micro-crystals were compressed, they either produced an electrical or magnetic field, or else produced changeable force fields that could raise the mid-section's temperature to a relatively high degree, thus causing the stored heat to flow in one direction. The resultant thrust of the air enabled the movement in any direction. The individual mini-crystals seemed to flutter rather than fly, and were incapable of steering an exact course—at least during the experiments conducted by the scientists in the laboratory. However, if they joined each other by chain-linking their wing tips, the ensuing aggregates possessed improved aerodynamic properties which increased proportionately with the number of links.

"Each crystal combined with three other crystals. In addition, its arm could link up with another crystal's middle section. This permitted a multilayered structure of ever-larger systems. The individual crystals did not even need to touch directly. It sufficed for the wing tips to come into close proximity to bring about a magnetic field which kept the entire system in balance. When a given quantity

> of "insects" clumped together, the aggregates then displayed definite and observable behaviour patterns. If the aggregate was subjected to external stimuli, it could change its direction, form, shape or the frequency of its internal impulses. Following such a change, the field would reverse its polarity, and as a result, the crystals no longer attracted but repelled each other and then broke down into their individual components."

It is on purpose that we give such a long quotation from Stanislaw Lem's story *The Invincible*. On the planet Regis III, humans came across an unusual phenomenon. Tiny primitive crystals having a primitive behaviour pattern under certain conditions came together into a cloud which acted as a great superorganism. The cloud possessed nearly inexhaustible adaptation abilities because it had an enormous memory consisting of the micro-memories of the crystals and could store a colossal amount of information.

Do you still regard this method of building a complex system out of elementary components as something extraordinary? We believe that after you have read the preceding chapters of this book, this method of organizing a complex behaviour should seem familiar to you. The observation of living organisms drives us to conclusions

which do not at all run counter to Lem's idea. Such aggregation of simple organisms into a much more complex organism is one of the ways in which living creatures evolved. K. Zavadsky, who has studied evolution for many years, called this method *synthesogenesis*. The period when unicellular algae turned into multicellular organisms was a decisive step in the progress of the organic world; an association of bees in a beehive or ants in an anthill are examples of the same phenomenon.

A mere concentration of identical subsystems or organisms, however, is not a new system or organism. A multitude of bees which meet in a summer meadow and belong to different families is not the same as an association of bees from the same beehive. A collection of passengers who happened to be travelling in the same train is quite different from a collection of buyers and sellers in a marketplace.

What is the big difference? In general, we can say that a collection of components may be called a unified system if these components have the potential of forming static or dynamic structures which are needed for the components and the whole collection to "survive", i.e. they must be able to interact to achieve local and global goals. This is no definition, certainly, but rather an attempt to find an answer to an extremely difficult question. A comprehensive investigation of this problem lies far beyond

our ability. Yet we are inclined to think that these reflections go into the nature of all the models of collective behaviour and interaction. Note also that when we discuss biotic communities we should bear in mind that in reality these potential properties are only realized in part while those not realized await an appropriate time. For example, some well-known experiments were carried out on a type of bacterium which always inhabited environments devoid of certain types of carbohydrate. When they were placed in an environment in which these "inedible" carbohydrates were the only food the bacteria could use, they started producing an enzyme to break them down. This ability was stored in their genetic structure for a "rainy day" and was realized when the need arose. Another example is the great potential abilities of all human beings, which for the most part are never realized in an individual or, possibly, in a social community.

Thus, synthesogenesis is a method for increasing the number of potential properties which may be handy in an emergency situation or unfamiliar environment.

Consider a simple model illustrating the potential of synthesogenesis. Figure 6.1 shows a torus world, i.e. a collection of cells placed on an external surface of the torus shaped like a doughnut. We assume that inside these cells there is food which may be used by "organisms" inhabiting

these cells. The role of the "organisms" in our example will be played by automata with linear tactics. In its simplest form, such an automaton is the one-action automaton given in Fig. 6.2*a*. In state *1* it re-

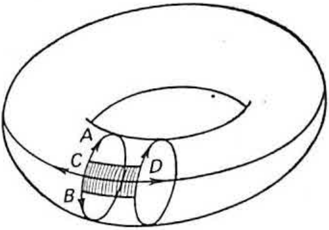

FIG. 6.1

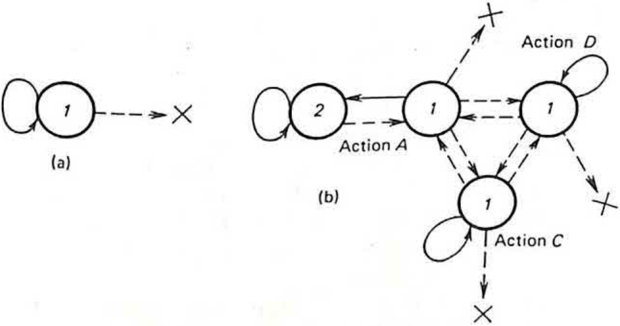

FIG. 6.2

ceives a fine signal and dies (in the figure this action is marked by a cross). The action which can be made by the automaton is a movement in a certain direction to the

next cell in the torus. We denote the four possible directions of movement A, B, C, D (see Fig. 6.1). Then the simplest automata will fall into four corresponding types designated by the same letters. We assume that the automata in one cell may combine. If this is a combination between two automata of the same type, the result is an increase in the length of the petal, i.e. in the memory capacity for this action. When different automata are combined, the newly-produced automaton has two petals instead of one. Figure 6.2*b* shows an automaton which was created by the aggregate of four automata: two belonging to the A type, one of the C type, and one of the D type. For convenience, we will designate such an automaton A^2DC.

Unlike a classical automaton with linear tactics, our automaton cannot accumulate fines infinitely and "dies" when the series of fines (shown by the broken lines) exceeds the number of the automaton's states (in the automaton given in Fig. 6.2*b* this number is four). Besides, the change of petals occurs equiprobably.

The fine and reward signals are generated by the environment in the following way. If an automaton in a cell eats all the food, it is rewarded; otherwise, it is fined. When the automaton is fed up (this takes it one time unit in this model) and leaves the cell, the cell is refilled with food immediately or remains empty until it is refilled

according to some law characterizing the environment.

If several automata find themselves in one cell, they are forced to combine to form a new and more complex organism.

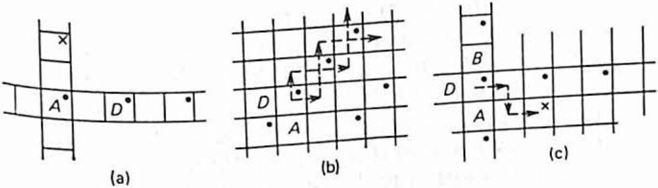

FIG. 6.3

Consider some moments of the evolution occurring on the torus.

Figure 6.3 gives a number of the simplest situations in a certain part of the toroidal surface. Cells containing food are marked with dots. It is assumed that a cell is fully refilled with food as soon as an automaton leaves it. Figure 6.3*a* shows two most elementary automata. Automaton *A* consumes the food in its cell and moves upward. In this circle, however, there is no food to be found. As a result, it starves to death in the cell marked with a cross. Automaton *D* is luckier. If there is food throughout the circle, it will move counterclockwise around the closed circle and never get hungry. It will live a life of ease for ever.

Figure 6.2*b* features another very simple situation. Automata *A* and *D* meet in a

মনের জানালা

অমিত চক্রবর্তী

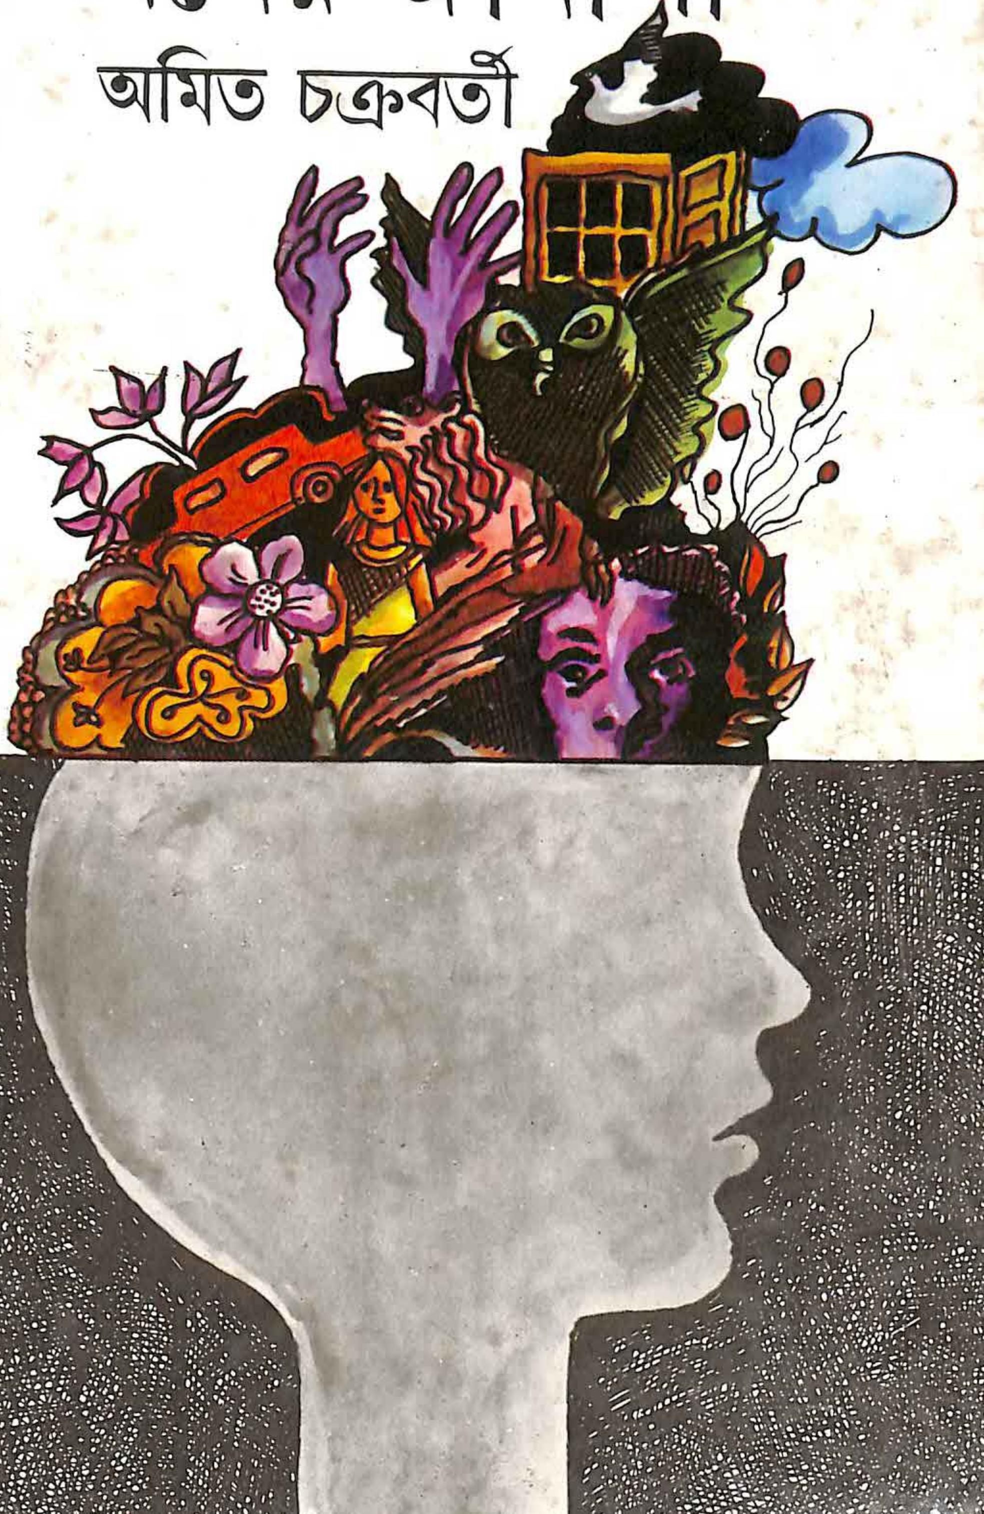

automata combine mechanically like the automata in our evolution model on the torus, the number of states grows as n^2 if each of them has n states. When they combine by random pair interaction, however, it allows them to function like automata with a memory capacity of 2^n. In Chap. 5, we also came across the "polymerization" phenomenon. When an automaton with 8 states joined the firing squad, it was able to rely on the memory of the whole automata chain without changing its own structure. This phenomenon seems extremely significant to us.

Alongside this process, another important process unfolds in biological and technological evolution, which shows that inhomogeneity in an organism starts growing as soon as specialized systems are introduced.

6.2. Segregatiogenesis and Its Effects

Like "synthesogenesis", the term "segregatiogenesis" was coined by K. Zavadsky. This means that during the evolution of biological species their complication through the agglomeration of simpler organisms into more complex ones is accompanied by a differentiation in the functions carried out by individual subsystems which results in structural changes in these subsystems so that they may perform better. Progress cannot be achieved by relying on multipurpose identical elements, which

are a sort of Jack of all trades. If the food on the torus in our example is arranged so that it would feed an automaton that moves like a knight in chess, the automaton's specific function must permit it to make this move in one time unit. If the food were differently arranged, this action would not be necessary at all.

A conflict between universality and specificity, homogeneity and heterogeneity is a general phenomenon inherent in any system including biocenoses and technocenoses.

The number of bees in a hive may vary widely. They each make up subsystems which can live independently. The queen bee of the hive, however, is the only one in the hive and it can not live long without its "subjects". Here differentiation has gone very far and one of the subsystems has already lost its stability to function independently of the system to which it belongs.

It is obvious, however, that such subsystems are useful. In Chap. 3, we showed how important inhomogeneity is in an automata collective. Reflex levels, and optimism-pessimism levels were the first indications of differences which, as they become more deep-rooted, allow a inhomogeneous collective to do a better job than a homogeneous one. We must admit that any automaton in this collective could also function on its own. This only means that specialization has not yet reached the limit

at which the independent existence of an individual subsystem becomes impossible. Thus, specialization is absolutely indispensable for progress for it is the only way a variety of goals can be achieved in a less time-consuming manner.

To illustrate this, we are going to trace the evolution of computers. At an early stage each computer was an indivisible single entity. Its processor, memory, peripheral devices, and control system were so rigidly connected to each other that they could not be even analyzed separately to say nothing of functioning on their own. All the processes occurring in a computer are strictly sequential and are governed by a central processing unit. This computer may be compared with a kind of a "cell" in a computer world.

How did computers develop in the following years? One of the major trends in computer evolution was associated with the increasing complication through introduction of new subsystems which permitted a wider number of functions*. Thus, graph plotters made it possible to generate graphical information rather than just the traditional alphanumeric format, while random number generators allow us to use methods that rely on random distributions. Although they make the "cell" more com-

* These subsystems may also be realized as software.

plex, new subsystems do not at all alter the underlying structure of the system. The complication, however, puts a great burden on the control system because it is assigned additional tasks. The moment came when the operating systems, which act as central controller watching all the processes occurring in a computer, became a source of trouble. It became clear that any further complication in computer structure hindered computer evolution. Computer scientists were discussing situations called deadlocks with increasing frequency. A deadlock is a state in which different processes occurring in a computer make conflicting demands and the machine does not know what to do.

It was obvious that centralized control could no longer permit further complication in computer structure and improve its efficiency.

The next inevitable step in computer evolution was the creation of computer systems or the grouping of several computers. That's when synthesogenesis came into action. "Unicellular" computers gave way to the next, "multicellular" generation. These groupings had a different structure. Figure 6.4 gives several types of multiple computer complexes. Figure 6.4*a* shows a structure with the central computer *1*, which acts as a central controller for computers *2*, *3*, *4*; in Fig. 6.4*b* you see a mixed structure and in Fig. 6.4*c* you see a decentralized

structure in which all computers involved enjoy equal rights. Note that even in the centralized structure there are traces of decentralization. The central computer does not closely watch the processes occurring in subordinate computers. It only initiates

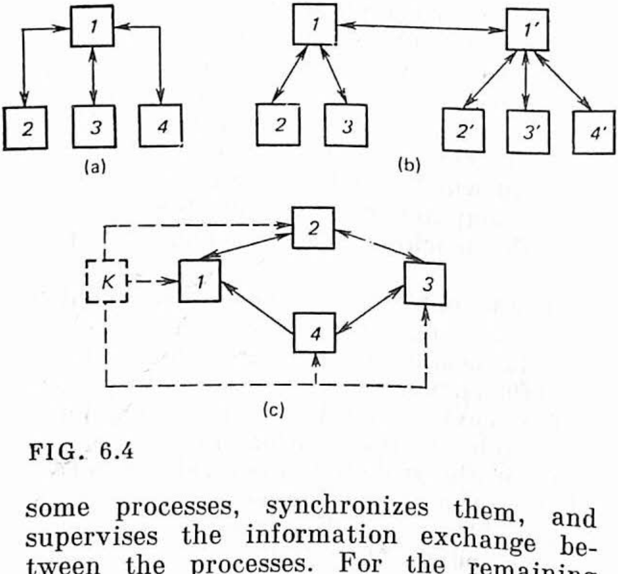

FIG. 6.4

some processes, synchronizes them, and supervises the information exchange between the processes. For the remaining functions, the machines integral to the complex operate independently. This trend in computer evolution shows that we were quite right not to follow yogis mentioned in Chap. 5.

It is noteworthy that the decentralized structure in Fig. 6.4*c* demonstrates how

a non-specific centralized control can be introduced into the structure of the "organism". A synchronization unit K, shown in the figure by a broken line, may transmit a signal to all the computers in the complex at one time over a common bus. For example, a signal interrupting all the calculations in order to receive new external information, or to repeat a calculation, or to verify data. A multiple-computer complex, however, may do without the central controller. Then a decentralized system will be synchronized in a fashion similar to the firing squad discussed in Chap. 5.

In parallel to this major trend in computer evolution, there was another method in which computers would be created on the basis of the homogeneous cellular structures also described in Chap. 5. This method is synthesogenesis in its purest form. It was believed that individual homogeneous and universal subsystems, i.e. automata in cells of homogeneous structure with potentially identical links between them, would allow a radical increase in computer performance. The idea, however, was fruitless because segregatiogenesis proved far more efficient.

At the next stage of computer evolution, an attempt was made to group very specialized subsystems and define the functions realized by each of them. Initially, the resulting structures resembled the type shown

in Fig. 6.4. The only difference was that the computers in the complex were specialized. Thus, they could be specially designed to process symbolic information, work with arrays, or do the preliminary processing of the information coming from an object of control. At the same time, like in the automata models with reflex or pessimism-optimism levels, they could also operate autonomously, i.e. outside the complex.

The process of segregatiogenesis led to the disappearance of this ability. The differentiation involved even the initial cell, i.e. the computer compared with a cell. Its integral components have, in a way, become independent. The result was a structure shown in Fig. 6.5. Here the processors, memories, exchange and control units seem to float in a computational environment. Their agglomeration into a structure occurs dynamically. Having received a task, the controllers look for idle performers and organize the process. In Fig. 6.5, you can see a moment when controller C_1 having been given a job, combines two processors P_1 and P_4, one memory M_2 and three exchange units E_2, E_3 and E_5. At the same time, controller C_2 is organizing another process by creating a "task force" which consists of processor P_3, memory M_1 and exchange unit E_1. In this arrangement, controllers receive computational problems from the external environment. Similarly,

exchange units receive initial data from the outside world too and the results are returned to the external environment. Having done their jobs, the structures disintegrate.

In this system, segregatiogenesis went as far as to deny individual subsystems an

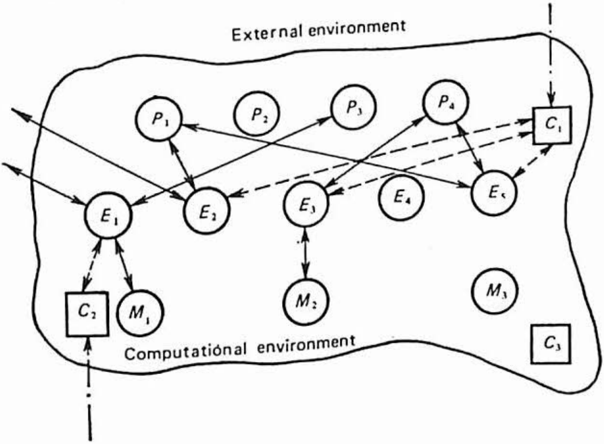

FIG. 6.5

opportunity to exist autonomously. This "organism" may only function as a structure with at least one controller, one exchange unit connected with processors or memories. The ability to establish structures that match the problem to be solved demonstrates the adaptability while specialization of the individual subsystems makes it possible to realize their inherent

function by concurrent and high speed operations.

A research into progress in evolution allowed K. Zavadsky to suggest the diagram given in Fig. 6.6. He holds that there are

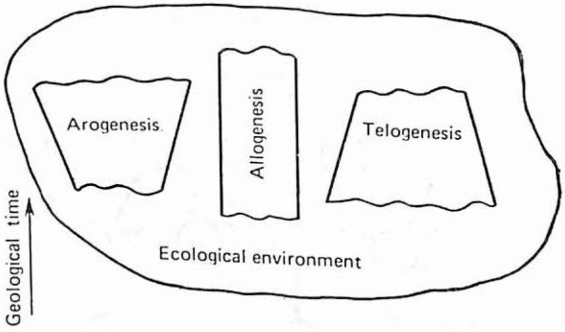

FIG. 6.6

three types of evolutionary development of biological organisms. The first is called *arogenesis* and implies the improvement of an organism's adaptability. The variety of environments in which it can survive and give offspring becomes wider. This process may unfold either through synthesogenesis, like our model of evolution on the surface of a torus, or through segregatiogenesis, like in a computational environment, which permits the solution of any problem within the range of the system's resources. Note that, for the same resources, an integrated system of the type

shown in Fig. 6.4*a* would not be able to cope with the organization of two concurrent processes as shown in Fig. 6.5.

Unlike arogenesis which means a system's better adaptability, *allogenesis* implies a change in functions realized by an organism to get a new, ecologically equivalent, function. In other words, allogenesis means the change of one ecological niche for another one which is better for the organism's survivability. Besides biology, this phenomenon also occurs in technology. Electronic calculators were used in design offices before the advent of computers. Later, when they were swept away by computers, calculators found another ecological niche in accounts offices where computers were not cost-efficient. When aircraft appeared, airships became practically extinct; today, however, enthusiastic admirers of airships seem to have found another ecological niche for them in the modern technocenosis and it is likely that we will soon see their unique outlines in the sky again.

Finally, *telogenesis* is like the other side of arogenesis. In telogenesis, there occurs a mere specific adaptation to a particular ecological environment by becoming highly specialized. The examples of telogenesis in technological systems are obvious. Nearly all specialized systems may be considered from this viewpoint. The primeval axe which could be used for a variety of purposes gave birth to a great number of cutting

tools, many of which can only be used for a particular job (like, for instance, a wood chopper unless you use its butt to drive in nails).

Arogenesis, allogenesis, and telogenesis are not alternative ways of evolution. They act at the same time and produce a combined effect. The domination of one of them in the development of some organism may only be temporary and passing. All three trends, however, help achieve the same goal, i.e. improve the organism's adaptability and, consequently, its survivability in a particular environment. The same may be said about technological systems.

6.3. Evolution in the Erehwon City

The strange name of this city is the word "nowhere" written backwards. It was coined by the English writer Samuel Butler in the second half of the 19th century for his novel of the same title. Butler's novel is written in the tradition of Utopian fiction. The main character, a young man named Higgs, travels in the mountains and finds himself in a most fantastic city. Erehwonians live according to laws which are very different from the morals and conventions of their European contemporaries. Thus, they consider disease and disaster to be crimes and those who are guilty are brought to court and punished. Nobody is cheered by the birth of a child, and children are never grate-

ful to their parents when they grow up. However, they have produced a race of great physical beauty and strength. Erehwonians first welcome Higgs but before long send him to prison.

The reason for this unexpected turn of Fortune's wheel is brought about by Higg's possession of a watch. He learns why the watch has frightened the Erehwonians so much from Yram, the daughter of the superintendent of the local prison. This reason is directly connected with the subject of this book.

Before we get down to discussing it in detail, let us turn to the life and art of Samuel Butler. A man of many talents and passions, he lived an eventful life. Among other things, one of his major ambitions was to comprehend the nature of evolution. Butler was greatly influenced by Charles Darwin and his theory of the origin and perpetuation of species. First a devoted Darwinian, he later doubted the validity of the theory. According to Butler, the most vulnerable point in Darwinism was that biological evolution was only a chain of random interactions and mutations. He was convinced that the process was purposeful*. But who could it be governed by? Butler was a man who accepted reason as

* The idea of purposefulness and rationalism of evolution was also supported by Academician L. Berg, the founder of the theory of nomogenesis.

the ultimate authority, and in his writings he criticized religion and more than once ridiculed religious practice and dogma. However, in books on evolution (*Life and Habit*, 1877; *Evolution, Old and New*, 1879; *Unconscious Memory*, 1880; *Luck or Cunning*, 1886), Samuel Butler denied Darwin's idea of the probabilistic nature of evolution. One of his arguments was his own concept of evolution in technology. Butler described this concept, probably for the first time in his article *Darwin Among the Machines*, which came out in 1863. He pointed out that in a technological evolution man was the link bringing purpose and rationalism into the evolutionary process. In *Erehwon*, he elaborated this idea.

Gradually, Higgs comes to realize that at one time Erehwon had a highly developed technocenosis which had been created by researchers and engineers to meet the needs of Erehwonians, facilitate their life, and stimulate further scientific and technological advance. Once created, however, the technocenosis started growing like a malignant tumor. Higgs happens to get hold of a copy of an Erehwon treatise *The Book of the Machines*. Thus he learns that technological advance swept over Erehwon at such a deadly pace that the citizens were steadily becoming slaves of the civilization of machines created by their own hands. The machines came to regard people as insects whose purpose was to

pollinate and fertilize mechanical devices, which could live their own independent lives. In criticizing the society he lived in, Samuel Butler wondered how many men live like the slaves of machines, serving them day and night from cradle to grave.

The same thing happened in Erehwon. The growing multitude of machines was perfectly adapted to a specially designed environment. They consumed energy produced for them and needed constant care. More and more citizens have to spend time on the machines: serving them, building new ones and finding jobs for them. It is hard to say what turn things could take in Erehwon existing in the author's imagination if he hadn't turned off the switch of technological progress in the city himself. A scientist appeared who used Darwin's theory to prove that before long Erehwonians would be completely conquered by the machines and that, as a result of segregatiogenesis, they wouldn't be able to exist independently. The result was the total destruction of machines and a ban on building any mechanism. Only remnants of the machines born during the time of technocenosis and kept in a museum reminded the citizens of the danger they managed to escape.

It is interesting to see what peculiarities, according to Butler, appeared in the course of evolution under man's will. First, the achievement of a goal by whatever means

available. The logic of a machine is different from that of man. Second, elaborate machines demand that the men controlling them be so highly specialized that they are only united by the data supplied by the system. We have already mentioned at the start of Chap. 5 how difficult it is to create general laws of control to supplement the logic of machines. As for the second peculiarity, Butler would be absolutely right if there were no chance in changing a human controller for a machine. This idea was discussed in Chaps. 3 and 4 and partially in Chap. 5. The image of a worker on an assembly line so brilliantly portrayed by Charlie Chaplin in *Modern Times* shows that Butler's fears were not groundless.

To a certain extent Butler was right. In 1978, R. Balandin wrote:

> "Even machines built and ruled by man have considerable power over us. Today we are as dependent on our machines as we are on the rest of nature. We eat technogenic, i.e. artificially selected and produced, species of animals and plants after their flesh has been technogenically treated. We live among machines and with the help of machines. It is obvious that we have to attend to them, take care of them and somehow try to adapt to them. We must take into account their capabilities and needs, often at the expense of our own interests. We must

promote further technological advance and make sure it is not marginal... ."*

We read this and are again haunted by the shadow of Erehwon.

Yet we cannot do what the Erehwonians chose to do. It is impossible to throw away the key to technological advance for none of us will ever be able to refuse the benefits it offers to man. What we need is a clear vision of the fact that under conditions in which the number of components in the different technocenoses is growing spontaneously, and in which global technocenoses embracing practically all human activities have appeared, control is of tremendous importance. The idea of control decentralization and introduction of integrated and coalition control systems is a major trend in exercising control over man-made supersystems.

6.4. Instead of a Conclusion. Evolution Goes on

Our narrative is coming to the end. We have sought to offer the general reader a popular account of the variety of methods that are used for building various decentralized control systems. We don't think that the similarities between biological, technological and administrative systems

* R. K. Balandin, *Geological Activity of Man. Technocenosis*, Visheishaya Shkola, Minsk, 1978.

are merely coincidental. Common control patterns are born from conditions that are common in the objects of control. The evolution of large-scale technological systems is one of the most convincing proofs of this statement. For this reason we are going to conclude by giving you another example of the evolutionary development of a large-scale technological system.

Computers appeared about forty years ago. In this chapter, we have already spoken about their evolution. The most recent breakthrough in the field is the introduction of a global data processing network, which is of a paramount importance both for computers and the mankind.

At an early stage of computer development, individual computers were connected by a cable. Later the use of the communication channels which already exist in the world's communication network was a qualitative leap forward. Before that a user who wanted to employ a computer to solve a problem knew exactly what computer would do the job and interacted directly with it. Now he doesn't know this. The problem may be solved by any of the computers integrated in the system. It is not infrequent that the distance between the user and the computer is rather long. Somehow the user has access to the network's total resource, which makes his capabilities infinitely greater. His problem may be solved by one computer in the network or by

several computers acting simultaneously. If you take into account the fact that both the computers integrated into the network and different parts of the problem to be solved may be inhomogeneous, you will easily see that organizing a problem's solution in this way can greatly improve efficiency.

The first territorial data processing network appeared in the USA in 1969. It was the ARPA network and it was followed by a generation of similar systems. It has a subnetwork composed of switching processors which ensures the exchange of information between all the computers in the network. In contrast to a telephone network, in which two subscribers are connected by a single channel, service requests in a data processing network are sent by the computers to switching processors. In a request, an addressee and a user are specified. Sometimes instead of an addressee the request contains the demands he should meet, i.e. a tolerable length of operation and the memory capacity required for the problem. If the addressee is available, the switching processor sends the request either directly (provided there is a direct communication channel between the processor and the addressee and the latter can take the order) or to another switching processor whose position in the network allows it to give the request to a proper addressee. If the addressee is not specified, the switching processor will itself deter-

mine which computer is appropriate for the job. Thus, a data processing network involves message switching rather than channel switching, with messages circulating around the network. Each message obtains a "transit visa" in each switching processor it encounters. Later these visas help the

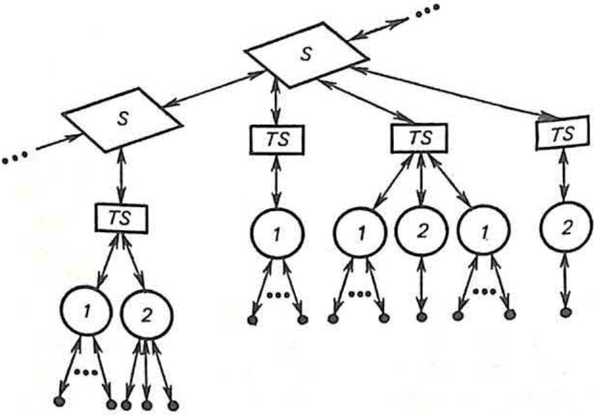

FIG. 6.7

switching processors send the messages back to the user who sent the message into the network.

For a switching processor to function easily the network also includes terminal switching processors, which serve as a kind of a buffer between the computer and the subnetwork of switching processors. Each terminal switching processor provides ac-

cess to a whole group of computers, each of which, in turn, uses its own terminals to serve dozens and hundreds of users. A fragment of such a network is given in Fig. 6.7. Here the data terminals are shown by black circles, the individual computers by big circles, the terminal switching processors *TS* by rectangles, and the switching processors *S* by parallelograms.

The ARPA network developed rapidly and before long it encompassed the whole of the US and soon reached Europe via communication centres in the UK and Norway. The next decade saw the advent of other networks. Thus, the TYMNET network which appeared in the USA allowed its users both to process information and to get an access to data banks storing an enormous amount of information on a variety of fields. Network designers believe that they will soon be able to offer the same services as libraries. A reader's request will come into the network and the target next will appear on a display panel. If there is a need to keep the text at home, the reader will get a hard copy from a printer.

The introduction of data processing networks has opened up new vistas in computer applications. Thus, the ARPA network was used to "publish a newspaper" on the problems of artificial intelligence and robotics. Correspondents wrote their articles on their personal computers and the articles were then collected in one comput-

er acting as a "newspaper". Each reader could obtain the "newspaper" and display it on his monitor, read it and, if necessary, print some or all of the text.

There was another unusual application of a data processing network when a conference on artificial intelligence was held in the USA. Both American and European scientists participated. The extraordinary thing was that neither the Europeans nor the Americans left home. All the reports were entered into the network so as to enable each participant to make his choice and have a better look at those reports which were most interesting, take part in a discussion or ask a reporter questions. The most pleasant thing was that anyone could have a break whenever he liked without running the risk of missing something really important. Because of the difference between the time zones of the USA and Europe, some participants were working, while others were sleeping so as to get back to work after a good night's rest.

Now back to the evolution of data processing networks. Besides American networks, there were other national territorial networks in a number of European countries. In 1974, the first phase of the CYCLADES network was put into operation in France. In 1971, eight European states signed a protocol concerning the development of a data processing network in Europe. In 1976, this network came into

operation with five switching centres in London, Zurich, Paris, Milan, and Ispra. This network was connected to the American networks through London, and centres were added in Vienna and many other cities. Some East European countries also established communication channels with the European network. Today there is a channel which connects Moscow to data processing in Europe. A great deal of efforts is underway to create a national Soviet network. Such is the rate of evolution of this new man-made system with its unprecedented power for data analysis and unforeseeable prospects for other computer applications.

Like any other system that develops as a result of an evolutionary process, the global data processing network does not have and cannot have an operational control centre. All the territorial national and international networks function autonomously. Coordinated control is achieved in a decentralized manner similar to the one which has been used in telephone networks for years. The service fees charged for the use of the computers, communication channels, and data banks are correlated with the cost of waiting time so as to raise or occasionally maximize the networks' efficiency. All the administrative information circulating between the switching centres are drawn up in the form of standardized data exchange protocols, which

make it easier for new segments and networks to join the main network.

So far synthesogenesis is obviously dominating in the world's data processing network. There are indications, however, that segregatiogenesis is starting to produce its effect. Some network members are beginning to specialize in handling problems of a particular type and equipment is being designed to meet these functions and thus be more profitable. It is likely that this trend will continue. The probable result is that separate segments of the network will become no longer autonomous and will not be able to "survive" independently.

In the future, the data processing network may fuse with the world television network and this will touch off a storm of fresh ideas and open up vistas that today can hardly be predicted. Man forges technological evolution with his own hands and unlike nature he does it purposefully. It is on man alone that the future of the technosphere depends.

There is more and more harmony in the music played by the orchestra of our man-made systems. It keeps on playing even though there is no conductor. For no conductor could ever handle so complex an orchestra.

To the Reader

Mir Publishers would be grateful for your comments on the content, translation and design of this book. We would also be pleased to receive any other suggestions you may wish to make.

Our address is:
Mir Publishers
2 Pervy Rizhsky Pereulok
I-110, GSP, Moscow, 129820
USSR

Also from Mir Publishers

Storming the Fortress of Fusion

by G. S. Voronov

The energy we get from the Sun is generated at its core in fusion reactions, in which hydrogen is turned into helium. If these reactions could be used on Earth, humanity would be able to provide itself with an abundance of energy. The problem of maintaining practical nuclear fusion has been one of the central problems of physics for some thirty years, and yet only now has it become clear that a solution is in fact possible. This book explains in a lively and informal way how the investigations have developed, the cleverness of the ideas used to tackle the problems, the ingenuity behind the experiments, and the successes and setbacks encountered in dealing with the capricious and cunning material we call plasma.
Intended for teachers, college students, and school pupils.

SCIENCE FOR EVERYONE

Paradoxes in the evolution of technological systems. Who designed the global telephone and communication network? The whole is the sum of the parts plus their interactions. Is a simple data register really simple? Automata exploring unknown worlds. Optimum or expediency? Random interactions and information exchange. Daydreaming and cynical automata. "I thought he thought I thought he. . . ." Heterogeneity in a group of automata is a key to success. Are we haunted by the ghost of Erehwon City? Puppets without strings: who is the puppeteer? Boundaries of centralization. How do systems evolve? A fresh look at design. Evolution goes on.

ISBN 5-03-000272-3